Landscape Analysis and Sample Complexity of Linear Quadratic Gaussian Control

Yang Zheng

Assistant Professor,
ECE Department, UC San Diego

Seminar at Department of Automation, Tsinghua University,
October 13, 2021

UC San Diego
JACOBS SCHOOL OF ENGINEERING
Electrical and Computer Engineering

SOC Lab

Acknowledgements

Yujie Tang
Harvard University

Na Li
Harvard University

Mayam Kamgarpour
EPFL

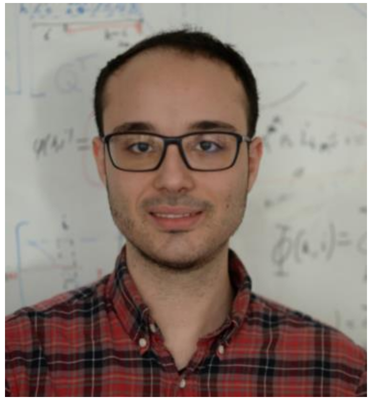

Luca Furieri
EPFL

Antonis Papachristodoulou
University of Oxford

Motivation

❑ Model-free methods and data-driven control

- Use direct policy updates;
- Become very popular in both academia and practice, from game playing, robotics, and drones, etc.

DeepMind

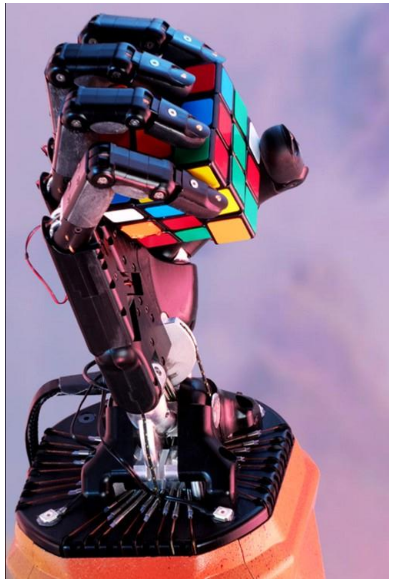

OpenAI

Applications

Duan et al. 2016; Silver et al., 2017; Dean et al., 2019; Tu and Recht, 2019;
Mania et al., 2019; Fazel et al., 2018; Recht, 2019;

Motivation

❑ Model-free methods and data-driven control

- but they are often computationally expensive, sample inefficient,
- lack non-asymptotic performance guarantees, such as sample complexity, safety, suboptimality, convergence etc.

Versions	Hardware
AlphaGo Fan	176 GPUs,[53] distributed
AlphaGo Lee	48 TPUs,[53] distributed
AlphaGo Master	4 TPUs,[53] single machine
AlphaGo Zero (40 block)	4 TPUs,[53] single machine
AlphaZero (20 block)	4 TPUs, single machine

Huge computation and
Millions of samples

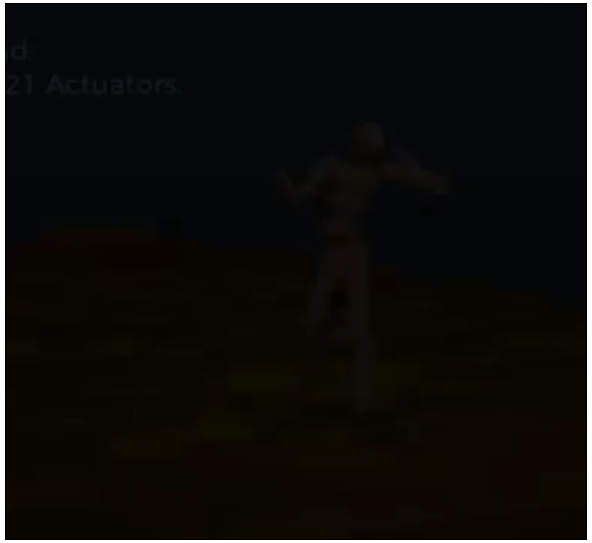

DeepMind

Uber running a red

- Highly nontrivial even for **linear dynamical systems!**

This talk

❑ Optimal Control

Feedback Paradigm

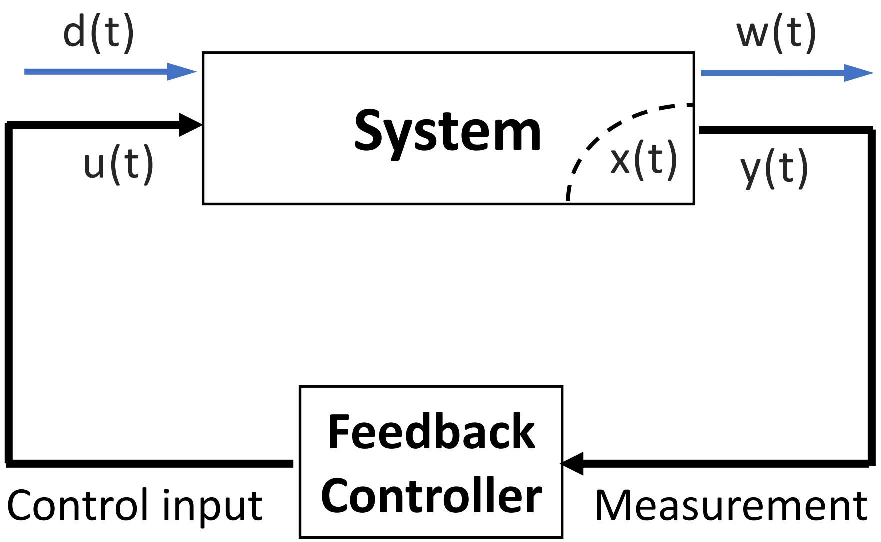

Control theory: the principled use of feedback loops and algorithms to drive a dynamical system to its desired goal

Linear Quadratic Optimal control

$$\min_{u_1, u_2, \ldots,} \quad \lim_{T \to \infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T}\left(x_t^{\mathsf{T}} Q x_t + u_t^{\mathsf{T}} R u_t\right)\right]$$

$$\text{subject to} \quad x_{t+1} = Ax_t + Bu_t + w_t$$

$$y_t = Cx_t + v_t$$

- Many practical applications
- Extensive classical results (Dynamic programming, Separation principle, Riccati equations, etc)
- **Linear Quadratic Regulator (LQR)** when the state x_t is directly observable
- **LQG** when only partial output y_t is observed

Major challenge: how to perform optimal control when the system is unknown?

Two main approaches

❑ Model-free: Direct policy iteration

- Give a parameterization of control policies; say **neural networks?** ✖
- Control theory already tells us many structural properties: **Linear feedback is sufficient for LQR**

$$u_t = Kx_t$$

$$\lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T}\left(x_t^{\mathsf{T}} Q x_t + u_t^{\mathsf{T}} R u_t\right)\right] := J(K)$$

Set of stabilizing controllers: $K \in \mathcal{K}$

LQR as an Optimization problem

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

Direct policy iteration

$$K_{i+1} = K_i - \alpha_i \nabla J(K_i)$$

✓ Good Landscape properties
  - Connected feasible region
  - Unique stationary point
  - Gradient dominance

✓ Fast global convergence (exponential)

A fast-growing list of references

- Fazel et al., ICML, 2018; Malik et al., 2019; Mohammadi et al., IEEE TAC, 2020; Li et al., 2019; Zhang et al., 2019; Furieri et al., 2019; Feiran Zhao & Keyou You, 2021, and many others

Two main approaches

❑ Model-based: Sys ID + robust control

- System ID + certainty equivalent control → adaptive control (Åström & Wittenmark, 2013).

Physical systems

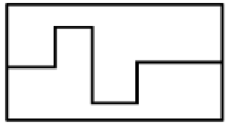

Control inputs u_t

Output meas. y_t

Estimated model

$\hat{A}, \hat{B}$

- Recent works → robust stability guarantees and sample complexity results, LQR problems (so-called system-level parameterization, Wang, Matni & Doyle, TAC, 2019)

Estimated model + uncertainty $\hat{A} + \Delta A, \quad \hat{B} + \Delta B, \qquad \|\Delta A\| \leq \epsilon_A, \|\Delta B\| \leq \epsilon_B,$

✓ Dean et al., 2020; Berberich et al., 2020; Boczar et al., 2018; Tsiamis et al., 2020; Umenberger et al., 2019; Yiwen Lu and Yilin Mo, 2021, and many others

❑ Library-based: Fundamental lemma (Coulson et al., 2019; Berberich et al., 2019; De Persis and Tesi, 2019)

Challenges for partially observed LQG

❑ Results on model-free or model-based LQG control are much fewer

- LQG is more sophisticated than LQR
- Requires dynamical controllers
- Its landscape properties are much richer and more complicated than LQR

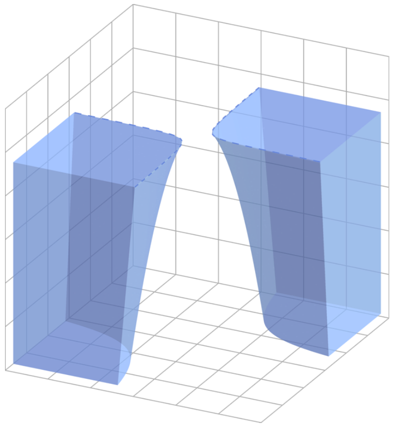

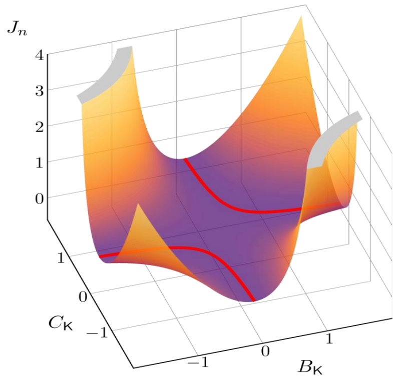

Part 1 Landscape Analysis

- The underlying technique, **system-level parameterization**, becomes non-trivial to use for the LQG case
- New techniques based on Input-output parameterization (IOP) (Furieri et al., 2019), are used for learning a robust LQG controller

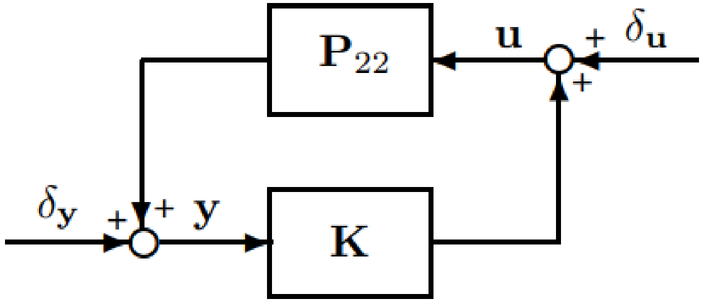

Part 2 Sample complexity

Today's talk

Part 1 Landscape Analysis

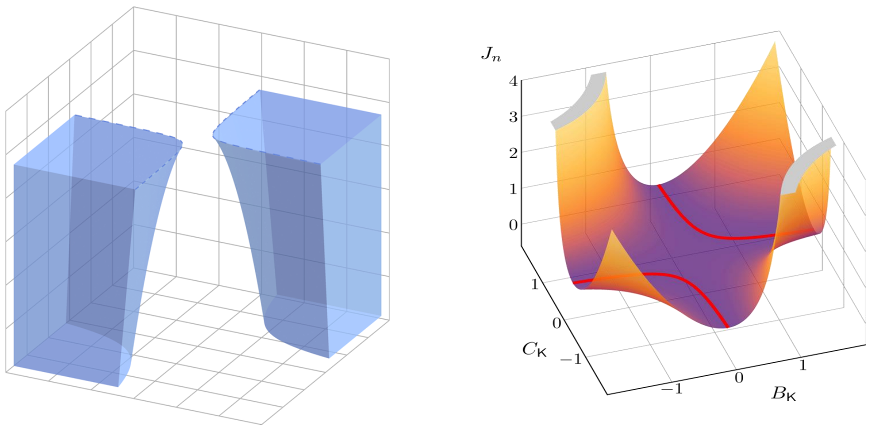

- Zheng, Yang, Yujie Tang, and Na Li. "Analysis of the Optimization Landscape of Linear Quadratic Gaussian (LQG) Control." arXiv preprint arXiv:2102.04393 (2021). link

Part 2: Sample Complexity

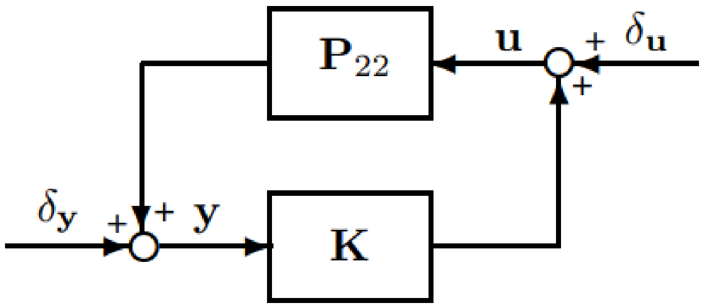

- Zheng, Y., Furieri, L., Kamgarpour, M., & Li, N. (2021, May). Sample complexity of linear quadratic gaussian (LQG) control for output feedback systems. In Learning for Dynamics and Control (pp. 559-570). PMLR. link

LQG Problem Setup

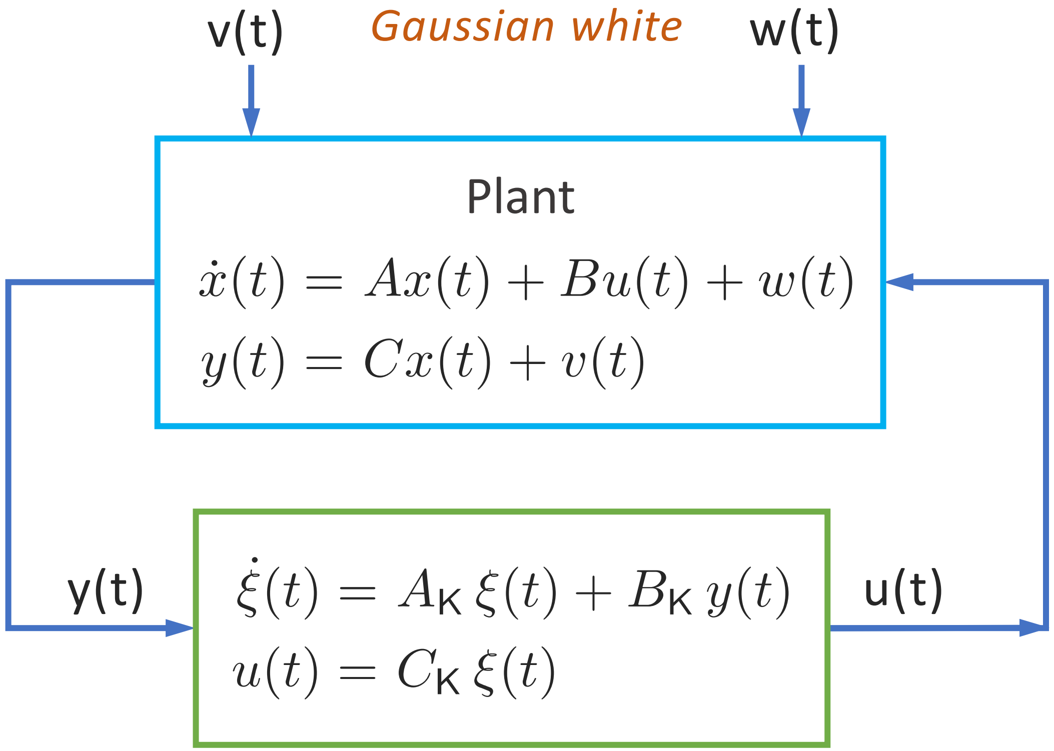

dynamical controller

$\mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}})$

Standard Assumption	(A, B), $(A, W^{1/2})$	Controllable
	(C, A), $(Q^{1/2}, A)$	Observable

Objective: The LQG cost

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

$\xi(t)$ 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.

* $(A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}})$ controllable and observable

Separation principle

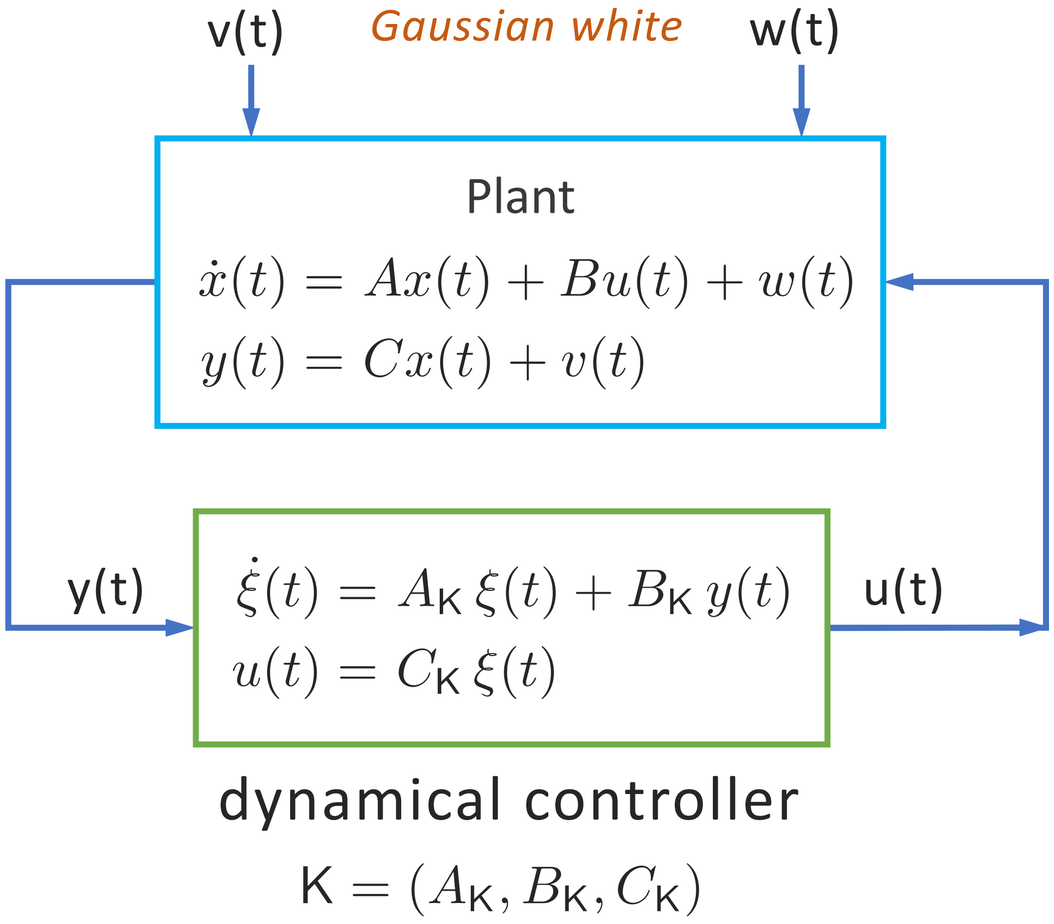

Explicit dependence on the dynamics

Objective: The LQG cost

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

Solution: Kalman filter + LQR based on the estimated state

$$\dot{\xi} = (A - BK)\xi + L(y - C\xi),$$
$$u = -K\xi.$$

Two Riccati equations

$$AP + PA^\mathsf{T} - PC^\mathsf{T}V^{-1}CP + W = 0,$$

Kalman gain $L = PC^\mathsf{T}V^{-1}$

$$A^\mathsf{T}S + SA - SBR^{-1}B^\mathsf{T}S + Q = 0$$

Feedback gain $K = R^{-1}B^\mathsf{T}S$

Model-free Optimization formulation

Closed-loop dynamics

$$\frac{d}{dt}\begin{bmatrix} x \\ \xi \end{bmatrix} = \begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix}\begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} I & 0 \\ 0 & B_{\mathsf{K}} \end{bmatrix}\begin{bmatrix} w \\ v \end{bmatrix}$$

$$\begin{bmatrix} y \\ u \end{bmatrix} = \begin{bmatrix} C & 0 \\ 0 & C_{\mathsf{K}} \end{bmatrix}\begin{bmatrix} x \\ \xi \end{bmatrix} + \begin{bmatrix} v \\ 0 \end{bmatrix}$$

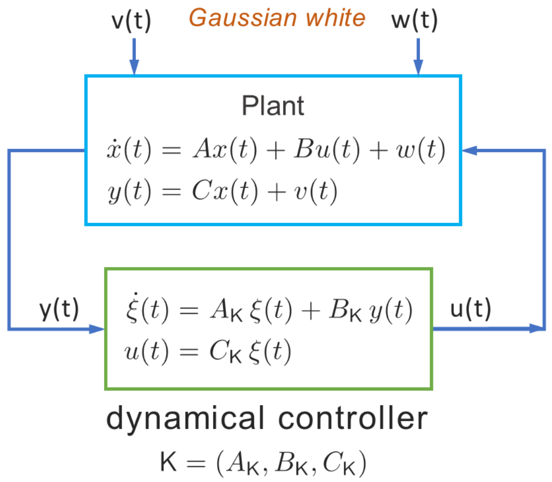

Feasible region of the controller parameters

$$\mathcal{C}_{\text{full}} = \left\{ \mathsf{K} \,\middle|\, \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \text{ is full-order, } \begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} \text{ is Hurwitz stable} \right\}$$

Cost function $\lim_{T\to+\infty} \frac{1}{T}\mathbb{E}\int_0^T (x^\top Qx + u^\top Ru)\,dt$

$$J(\mathsf{K}) = \operatorname{tr}\left(\begin{bmatrix} Q & 0 \\ 0 & C_{\mathsf{K}}^{\mathsf{T}}RC_{\mathsf{K}} \end{bmatrix} X_{\mathsf{K}}\right) = \operatorname{tr}\left(\begin{bmatrix} W & 0 \\ 0 & B_{\mathsf{K}}VB_{\mathsf{K}}^{\mathsf{T}} \end{bmatrix} Y_{\mathsf{K}}\right)$$

$X_{\mathsf{K}}, Y_{\mathsf{K}}$ Solution to Lyapunov equations

LQG as an Optimization problem

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

Direct policy iteration $\mathsf{K}_{i+1} = \mathsf{K}_i - \alpha_i \nabla J(\mathsf{K}_i)$

- ✓ **Does it converge at all?**
- ✓ **Converge to which point?**
- ✓ **Convergence speed?**

Landscape Analysis

Model-free Optimization formulation

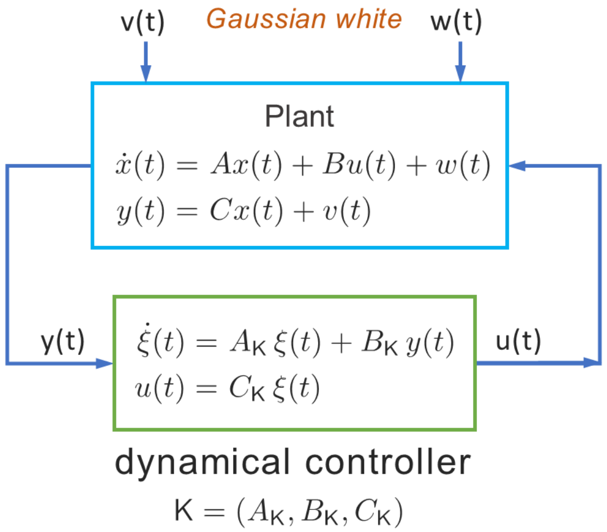

LQG as an Optimization problem

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

Landscape Analysis

- **Q1: Connectivity of the feasible region $\mathcal{C}_{\text{full}}$**
  - Is it connected?
  - If not, how many connected components can it have?
- **Q2: Structure of stationary points of $J(\mathsf{K})$**
  - Are there spurious (strictly suboptimal, saddle) stationary points?
  - How to check if a stationary point is globally optimal?

Connectivity of the feasible region

- **Simple observation: non-convex and unbounded**

Lemma 1: the set $\mathcal{C}_{\text{full}}$ is non-empty, unbounded, and can be non-convex.

Example:

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

$$\mathcal{C}_{\text{full}} = \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \,\middle|\, \begin{bmatrix} 1 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \text{ is stable} \right\}.$$

$$\mathsf{K}^{(1)} = \begin{bmatrix} 0 & 2 \\ -2 & -2 \end{bmatrix}, \qquad \mathsf{K}^{(2)} = \begin{bmatrix} 0 & -2 \\ 2 & -2 \end{bmatrix}$$

Stabilize the plant, and thus belong to $\mathcal{C}_{\text{full}}$

$$\hat{\mathsf{K}} = \frac{1}{2}\left(\mathsf{K}^{(1)} + \mathsf{K}^{(2)}\right) = \begin{bmatrix} 0 & 0 \\ 0 & -2 \end{bmatrix}$$

Fails to stabilize the plant, and thus outside $\mathcal{C}_{\text{full}}$

Connectivity of the feasible region

- **Main Result 1: dis-connectivity**

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

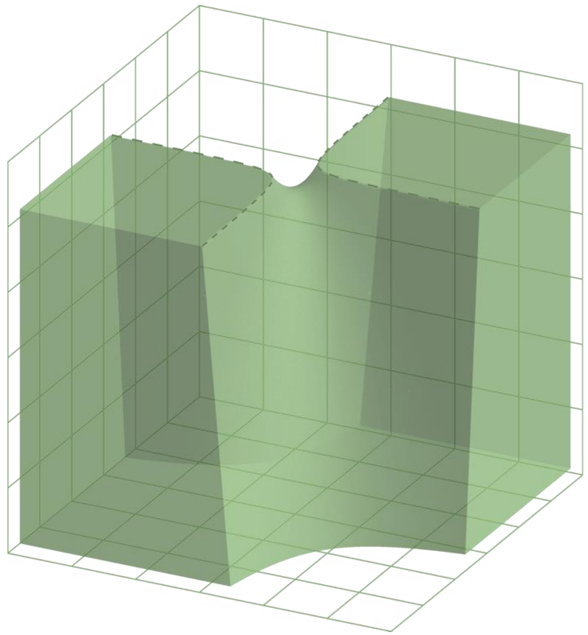

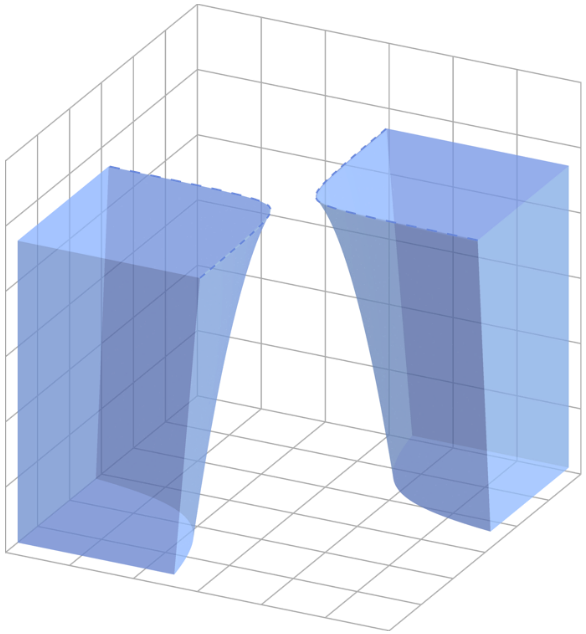

- ✓ Different from the connectivity of static stabilizing state-feedback controllers, which is always connected!
- ✓ Is this a negative result for gradient-based algorithms? → **No**

Connectivity of the feasible region

- **Main Result 2: dis-connectivity**

Theorem 2: If $\mathcal{C}_{\text{full}}$ has 2 connected components, then there is a smooth bijection T between the 2 connected components that has the same cost function value

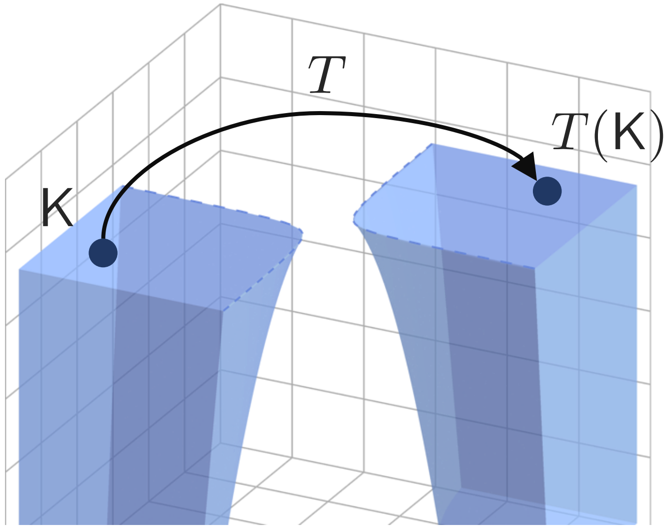

$$J(\mathsf{K}) = J(T(\mathsf{K}))$$

✓ In fact, the bijection T is defined by a similarity transformation (change of controller state coordinate)

$$\mathscr{T}_T(\mathsf{K}) := \begin{bmatrix} D_{\mathsf{K}} & C_{\mathsf{K}}T^{-1} \\ TB_{\mathsf{K}} & TA_{\mathsf{K}}T^{-1} \end{bmatrix}.$$

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

Connectivity of the feasible region

- **Main Result 3: conditions for connectivity**

Theorem 3: 1) $\mathcal{C}_{\text{full}}$ is connected if there exists a reduced-order stabilizing controller.

2) The sufficient condition above becomes necessary if the plant is single-input or single-output.

Corollary 1: Given any open-loop stable plant, the set of stabilizing controllers $\mathcal{C}_{\text{full}}$ is connected.

Example: Open-loop stable system

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

Routh--Hurwitz stability criterion

$$\mathcal{C}_{\text{full}} = \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < 1, B_{\mathsf{K}} C_{\mathsf{K}} < -A_{\mathsf{K}} \right\}.$$

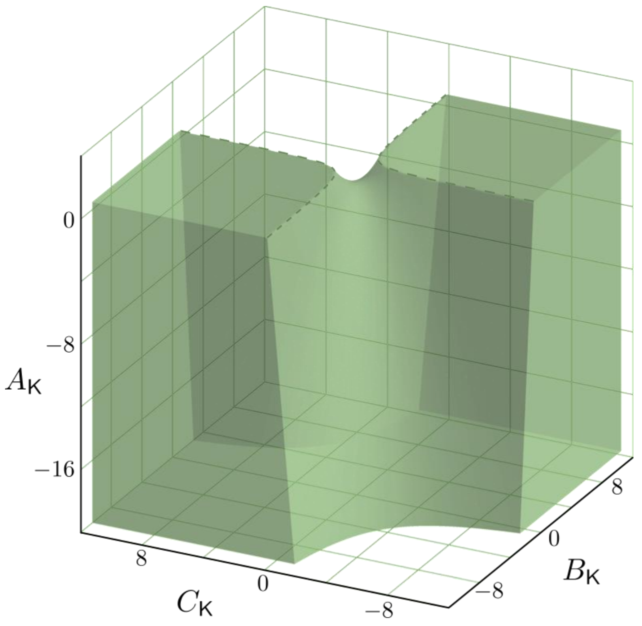

Connectivity of the feasible region

- **Main Result 3: conditions for connectivity**

Example: Open-loop unstable system (SISO)

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

- **Routh--Hurwitz stability criterion**

$$\mathcal{C}_{\text{full}} = \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| \begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} \text{ is stable} \right\}$$
$$= \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < -1,\ B_{\mathsf{K}}C_{\mathsf{K}} < A_{\mathsf{K}} \right\}.$$

- **Two path-connected components**

$$\mathcal{C}_1^+ := \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < -1,\ B_{\mathsf{K}}C_{\mathsf{K}} < A_{\mathsf{K}},\ B_{\mathsf{K}} > 0 \right\},$$
$$\mathcal{C}_1^- := \left\{ \mathsf{K} = \begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix} \in \mathbb{R}^{2\times 2} \middle| A_{\mathsf{K}} < -1,\ B_{\mathsf{K}}C_{\mathsf{K}} < A_{\mathsf{K}},\ B_{\mathsf{K}} < 0 \right\}.$$

Disconnected feasible region

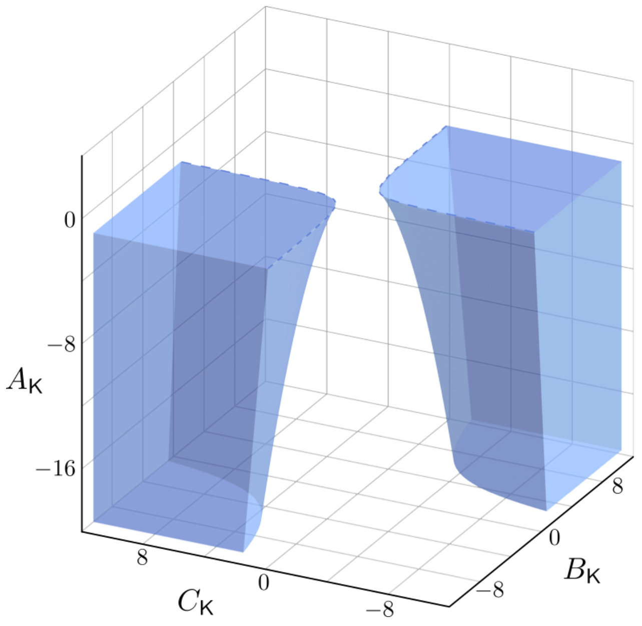

Proof idea: Lifting via Change of Variables

- **Change of variables in state-space domain: Lyapunov theory**

  - Connectivity of the static stabilizing state feedback gains

$$\{K \in \mathbb{R}^{m\times n} \mid A - BK \text{ is stable}\}$$

$$\iff \{K \in \mathbb{R}^{m\times n} \mid \exists P \succ 0, P(A-BK)^{\mathsf{T}} + (A-BK)P \prec 0\}$$

$$\iff \{K \in \mathbb{R}^{m\times n} \mid \exists P \succ 0, PA^{\mathsf{T}} - L^{\mathsf{T}}B^{\mathsf{T}} + AP - BL \prec 0, L = KP\}$$

$$\iff \{K = LP^{-1} \in \mathbb{R}^{m\times n} \mid \exists P \succ 0, PA^{\mathsf{T}} - L^{\mathsf{T}}B^{\mathsf{T}} + AP - BL \prec 0\}.$$

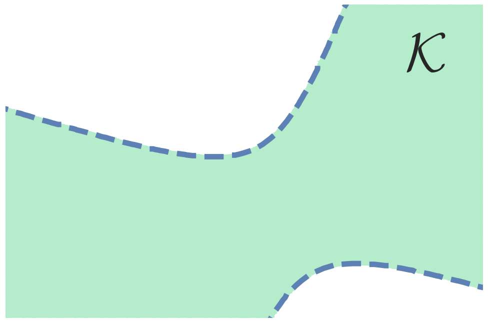

Open, connected, possibly nonconvex

  - How about the set of stabilizing dynamical controllers

$$\begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} \text{ is stable}$$

$$\iff \quad \exists P \succ 0,\ P\begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix}^{\mathsf{T}} + \begin{bmatrix} A & BC_{\mathsf{K}} \\ B_{\mathsf{K}}C & A_{\mathsf{K}} \end{bmatrix} P \prec 0,$$

Change of variables for output feedback control is highly non-trivial

[Scherer et al., IEEE TAC 1997]
[Gahinet and Apkarian, 1994]

Proof idea: Lifting via Change of Variables

- **Change of variables in state-space domain: Lyapunov theory**

[Scherer et al., IEEE TAC 1997]
[Gahinet and Apkarian, 1994]

$$\Phi(\mathsf{Z}) = \begin{bmatrix} \Phi_D(\mathsf{Z}) & \Phi_C(\mathsf{Z}) \\ \Phi_B(\mathsf{Z}) & \Phi_A(\mathsf{Z}) \end{bmatrix} := \begin{bmatrix} I & 0 \\ YB & \Xi \end{bmatrix}^{-1} \begin{bmatrix} G & H \\ F & M - YAX \end{bmatrix} \begin{bmatrix} I & CX \\ 0 & \Pi \end{bmatrix}^{-1}.$$

at most 2 connected components

$$\mathcal{F} \times \mathrm{GL} \xrightarrow[\text{surjective}]{\Phi} \mathcal{C}_{\text{full}}$$

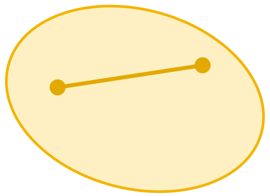

Convex thus connected

General linear group: the set of invertible matrices (similarity transformation)

2 connected components

$$\mathrm{GL}_n^+ = \{\Pi \in \mathbb{R}^{n\times n} \mid \det \Pi > 0\},$$
$$\mathrm{GL}_n^- = \{\Pi \in \mathbb{R}^{n\times n} \mid \det \Pi < 0\}.$$

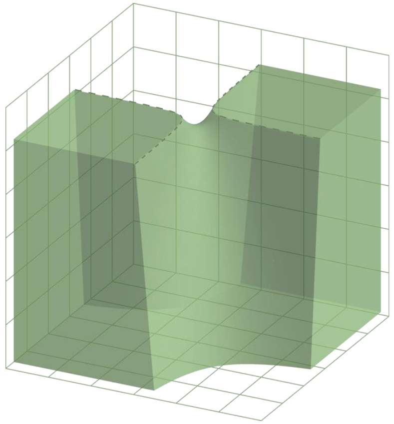

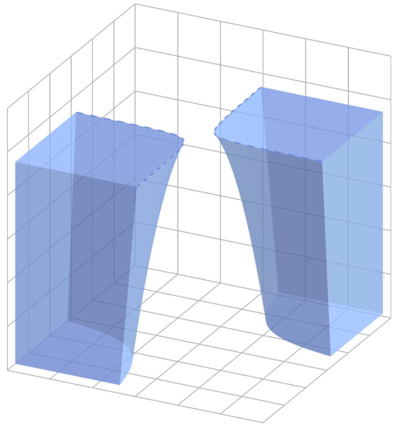

Model-free Optimization formulation

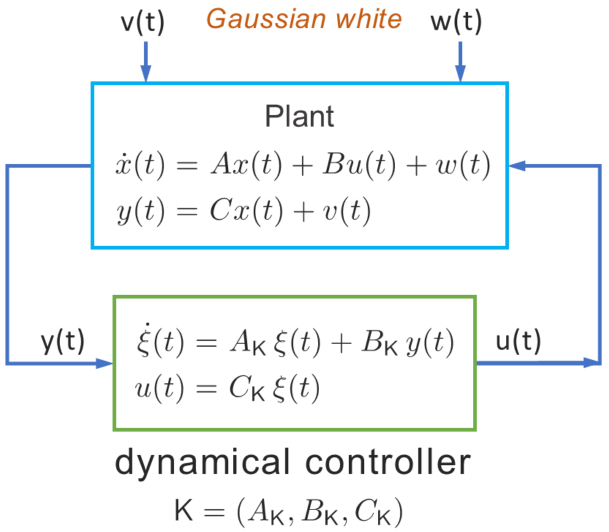

LQG as an Optimization problem

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

- **Q1: Connectivity of the feasible region $\mathcal{C}_{\text{full}}$**
  - Is it connected? **No**
  - How many connected components can it have? **Two**
- **Q2: Structure of stationary points of $J(\mathsf{K})$**
  - Are there spurious (strictly suboptimal, saddle) stationary points?
  - How to check if a stationary point is globally optimal?

Landscape Analysis

Structure of Stationary Points

LQG as an Optimization problem

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

- **Simple observations**

1) $J(\mathsf{K})$ is a real analytic function over its domain (smooth, infinitely differentiable)
2) $J(\mathsf{K})$ has **non-unique** and **non-isolated** global optima

Similarity transformation

$$(A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \mapsto (TA_{\mathsf{K}}T^{-1}, TB_{\mathsf{K}}, C_{\mathsf{K}}T^{-1})$$

$$\dot{\xi}(t) = A_{\mathsf{K}}\,\xi(t) + B_{\mathsf{K}}\,y(t)$$
$$u(t) = C_{\mathsf{K}}\,\xi(t)$$

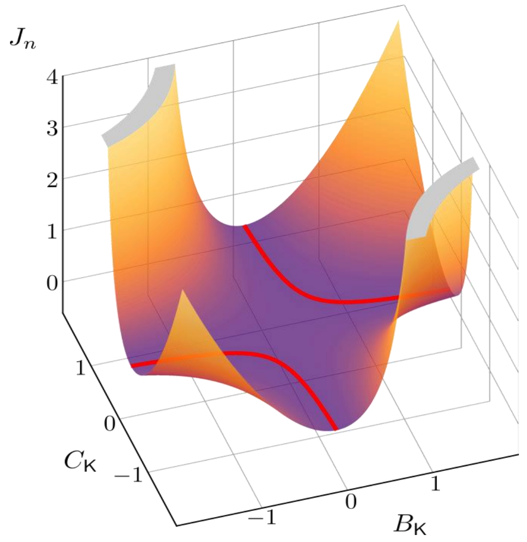

- $J(\mathsf{K})$ is invariant under similarity transformations.
- It has many stationary points, unlike the LQR with a unique stationary point

Structure of Stationary Points

- **Gradient computation**

Lemma 1: For every $\mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}}$, we have

$$
\begin{aligned}
\frac{\partial J(\mathsf{K})}{\partial A_{\mathsf{K}}} &= 2\left(Y_{12}^{\mathsf{T}} X_{12} + Y_{22} X_{22}\right), \\
\frac{\partial J(\mathsf{K})}{\partial B_{\mathsf{K}}} &= 2\left(Y_{22} B_{\mathsf{K}} V + Y_{22} X_{12}^{\mathsf{T}} C^{\mathsf{T}} + Y_{12}^{\mathsf{T}} X_{11} C^{\mathsf{T}}\right), \\
\frac{\partial J(\mathsf{K})}{\partial C_{\mathsf{K}}} &= 2\left(R C_{\mathsf{K}} X_{22} + B^{\mathsf{T}} Y_{11} X_{12} + B^{\mathsf{T}} Y_{12} X_{22}\right),
\end{aligned}
$$

where $X_{\mathsf{K}} = \begin{bmatrix} X_{11} & X_{12} \\ X_{12}^{\mathsf{T}} & X_{22} \end{bmatrix}, \qquad Y_{\mathsf{K}} = \begin{bmatrix} Y_{11} & Y_{12} \\ Y_{12}^{\mathsf{T}} & Y_{22} \end{bmatrix}$

are the unique solutions to two Lyapunov equations

LQG as an Optimization problem

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

How does the set of Stationary Points look like?

$$
\left\{ \mathsf{K} \in \mathcal{C}_{\text{full}} \;\middle|\; \frac{\partial J(\mathsf{K})}{\partial A_{\mathsf{K}}} = 0, \; \frac{\partial J(\mathsf{K})}{\partial B_{\mathsf{K}}} = 0, \; \frac{\partial J(\mathsf{K})}{\partial C_{\mathsf{K}}} = 0, \right\}
$$

- Non-unique, non-isolated
- Local minimum, local maximum, saddle points, or globally minimum?

Structure of Stationary Points

❑ Main Result

Theorem 4: Consider any open-loop stable plant. The zero controller with any stable A_{K}

$$\mathsf{K} = (A_{\mathsf{K}}, 0, 0) \in \mathcal{C}_{\text{full}}$$

is a stationary point. Furthermore, the corresponding hessian is either indefinite (strict saddle point) or equal to zero.

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

$$Q = 1, R = 1, V = 1, W = 1$$

Stationary point $\mathsf{K}^\star = \begin{bmatrix} 0 & 0 \\ 0 & a \end{bmatrix} \in \mathbb{R}^{2\times 2}, \quad \text{with } a < 0$

Cost function:
$$J\left(\begin{bmatrix} 0 & C_{\mathsf{K}} \\ B_{\mathsf{K}} & A_{\mathsf{K}} \end{bmatrix}\right) = \frac{A_{\mathsf{K}}^2 - A_{\mathsf{K}}(1 + B_{\mathsf{K}}^2 C_{\mathsf{K}}^2) - B_{\mathsf{K}} C_{\mathsf{K}}(1 - 3B_{\mathsf{K}} C_{\mathsf{K}} + B_{\mathsf{K}}^2 C_{\mathsf{K}}^2)}{2(-1 + A_{\mathsf{K}})(A_{\mathsf{K}} + B_{\mathsf{K}} C_{\mathsf{K}})}.$$

Hessian:
$$\left.\begin{bmatrix} \frac{\partial J^2(\mathsf{K})}{\partial A_{\mathsf{K}}^2} & \frac{\partial J^2(\mathsf{K})}{\partial A_{\mathsf{K}} \partial B_{\mathsf{K}}} & \frac{\partial J^2(\mathsf{K})}{\partial A_{\mathsf{K}} \partial C_{\mathsf{K}}} \\ \frac{\partial J^2(\mathsf{K})}{\partial B_{\mathsf{K}} A_{\mathsf{K}}} & \frac{\partial J^2(\mathsf{K})}{\partial B_{\mathsf{K}}^2} & \frac{\partial J^2(\mathsf{K})}{\partial B_{\mathsf{K}} \partial C_{\mathsf{K}}} \\ \frac{\partial J^2(\mathsf{K})}{\partial C_{\mathsf{K}} A_{\mathsf{K}}} & \frac{\partial J^2(\mathsf{K})}{\partial C_{\mathsf{K}} B_{\mathsf{K}}} & \frac{\partial J^2(\mathsf{K})}{\partial \partial C_{\mathsf{K}}^2} \end{bmatrix}\right|_{\mathsf{K}^\star = \begin{bmatrix} 0 & 0 \\ 0 & a \end{bmatrix}} = \frac{1}{2(1-a)} \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix},$$

Indefinite with eigenvalues:

0 and $\pm\frac{1}{2(1-a)}$

Structure of Stationary Points

❑ Main Result

Theorem 4: Consider any open-loop stable plant. The zero controller with any stable A_{K}

$$\mathsf{K} = (A_{\mathsf{K}}, 0, 0) \in \mathcal{C}_{\text{full}}$$

is a stationary point. Furthermore, the corresponding hessian is either indefinite (strict saddle point) or equal to zero.

Another example with zero Hessian

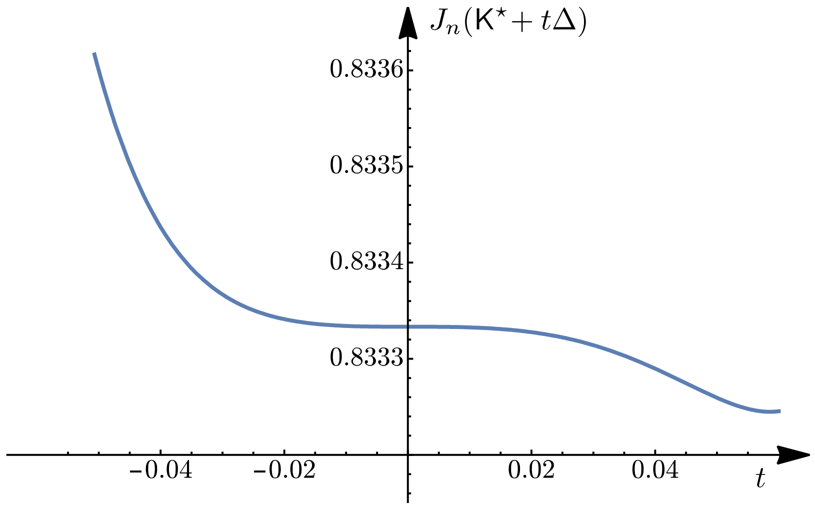

All bad stationary points correspond to non-minimal controllers

$$\left\{ \mathsf{K} \in \mathcal{C}_{\text{full}} \;\middle|\; \begin{aligned} \frac{\partial J(\mathsf{K})}{\partial A_{\mathsf{K}}} &= 0, \\ \frac{\partial J(\mathsf{K})}{\partial B_{\mathsf{K}}} &= 0, \\ \frac{\partial J(\mathsf{K})}{\partial C_{\mathsf{K}}} &= 0, \end{aligned} \right\}$$

Structure of Stationary Points

- **Main Result**

Theorem 5: All stationary points corresponding to controllable and observable controllers are globally minimal!!

$$\left\{ \mathsf{K} \in \mathcal{C}_{\text{full}} \;\middle|\; \frac{\partial J(\mathsf{K})}{\partial A_{\mathsf{K}}} = 0, \frac{\partial J(\mathsf{K})}{\partial B_{\mathsf{K}}} = 0, \frac{\partial J(\mathsf{K})}{\partial C_{\mathsf{K}}} = 0, \right\}$$

Particularly, given a stationary point that is a **minimal** controller

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. They are connected by a similarity transformation.

Example 1

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

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

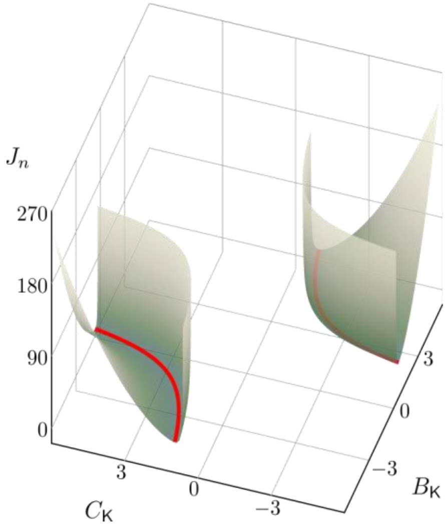

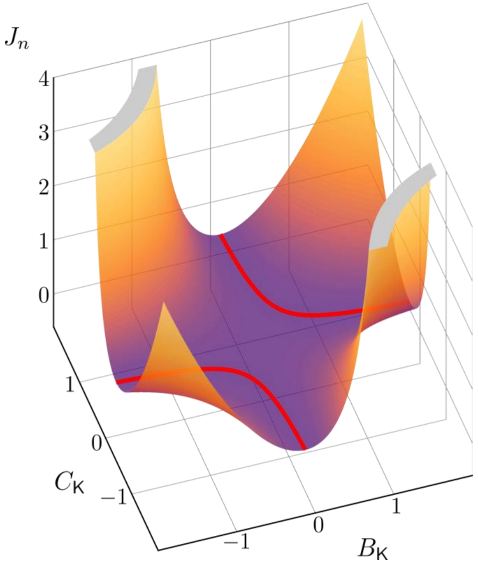

Example 2

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

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

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.

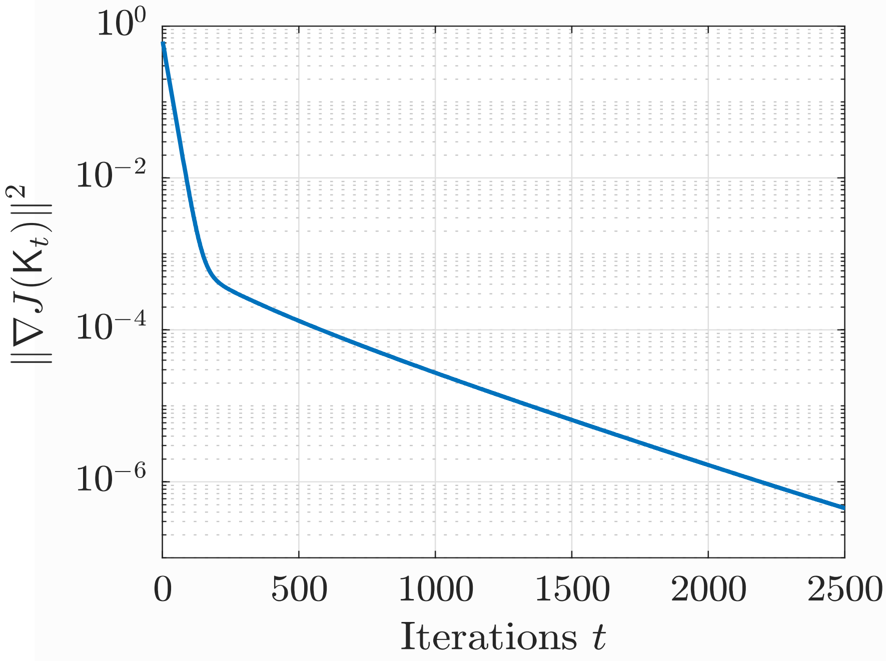

Open questions:

- ✓ Convergence conditions?
- ✓ Convergence speed?
- ✓ Alternative model-free parameterization

Comparison with LQR

	LQR as an Optimization problem $$\min_{K} \; J(K) \quad \text{s.t.} \; K \in \mathcal{K}$$	**LQG as an Optimization problem** $$\min_{\mathsf{K}} \; J(\mathsf{K}) \quad \text{s.t.} \; \mathsf{K} = (A_{\mathsf{K}}, B_{\mathsf{K}}, C_{\mathsf{K}}) \in \mathcal{C}_{\text{full}}$$
Connectivity of feasible region	❖ Always connected	❖ Disconnected, but at most 2 connected comp. ❖ They are almost identical to each other
Stationary points	❖ Unique	❖ Non-unique, non-isolated stationary points ❖ Spurious stationary points (saddle, nonminimal controller) ❖ **All mini. stationary points are globally optimal**
Gradient Descent	❖ Gradient dominance ❖ Global fast convergence (like strictly convex)	❖ No gradient dominance ❖ Local convergence/speed (**unknown**) ❖ **Many open questions**
References	**Fazel et al., ICML, 2018**; Malik et al., 2019; Mohammadi et al., IEEE TAC, 2020; Li et al., 2019; Zhang et al., 2019; Furieri et al., 2019; Feiran Zhao & Keyou You, 2021, and many others	Zheng, Tang, Li. 2021, link

Today's talk

Part 1 Landscape Analysis

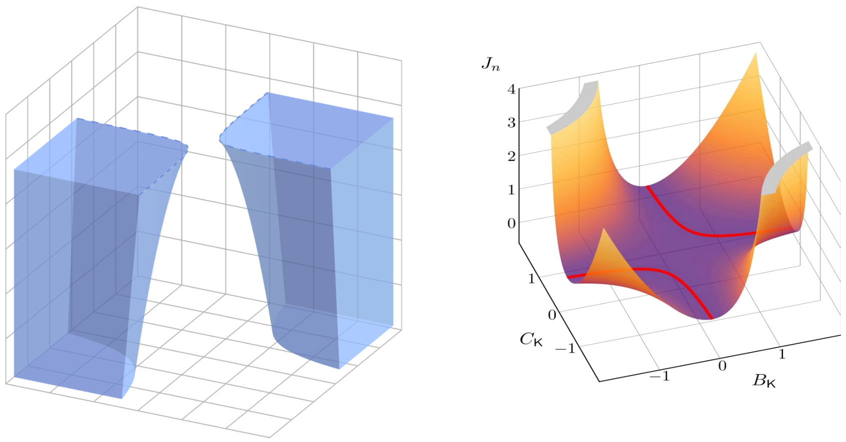

- Zheng, Yang, Yujie Tang, and Na Li. "Analysis of the Optimization Landscape of Linear Quadratic Gaussian (LQG) Control." arXiv preprint arXiv:2102.04393 (2021). link

Part 2: Sample Complexity

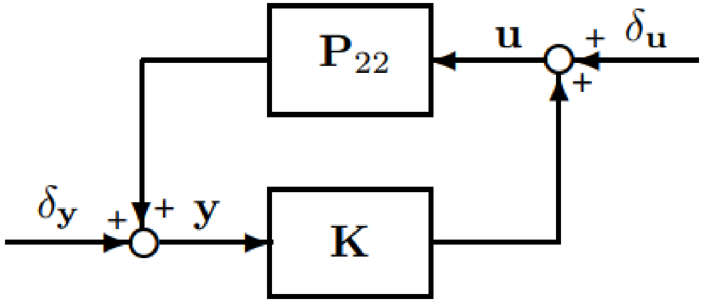

- Zheng, Y., Furieri, L., Kamgarpour, M., & Li, N. (2021, May). Sample complexity of linear quadratic gaussian (LQG) control for output feedback systems. In Learning for Dynamics and Control (pp. 559-570). PMLR. link

System ID + Robust Control

Physical systems

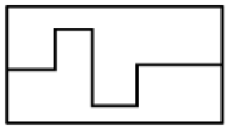

Control inputs u_t

Output meas. y_t

Estimated a dynamical model

- **How to represent a dynamical system: space-space or frequency domain?**
  - ✓ State-feedback LQR seems easier

$$\hat{A} + \Delta A, \quad \hat{B} + \Delta B, \qquad \|\Delta A\| \leq \epsilon_A, \|\Delta B\| \leq \epsilon_B,$$

  - ✓ Then use a recent tool called system-level parameterization (SLP, frequency domain technique) for robust control and sample complexity analysis; see Dean et al., 2020

- **Partially observed LQG case**

Natural idea: estimate $\|\hat{A} - A_\star\|, \quad \|\hat{B} - B_\star\|, \quad \|\hat{C} - C_\star\|,$

Then, design a robust LQG controller?

Highly Non-trivial

- ✓ Dean et al. 2020 works only for state feedback via SLP
- ✓ The realization of A, B, C is not unique!!

Frequency domain formulation

- **State-space model**

$$x_{t+1} = A_\star x_t + B_\star u_t + B_\star w_t, \quad y_t = C_\star x_t + v_t \,.$$

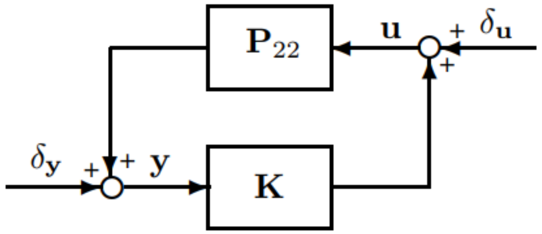

Least-square fits a coarse model

- **Unique transfer function**

$$\mathbf{G}_\star(z) = C_\star(zI - A_\star)^{-1}B_\star \,,$$

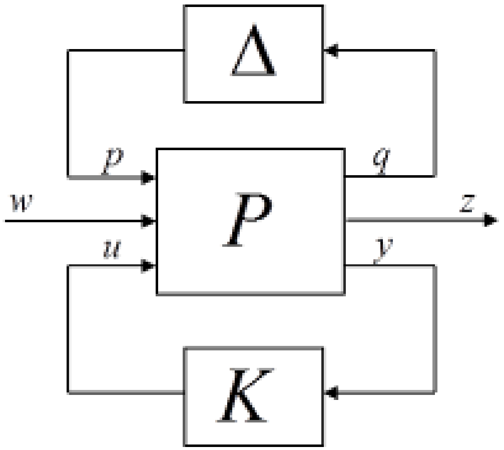

High dimen. stats bounds the error

Estimate a nominal model as well as its uncertainty

$$\|\boldsymbol{\Delta}\|_\infty := \|\mathbf{G}_\star - \hat{\mathbf{G}}\|_\infty < \epsilon$$

Design a robust LQG controller

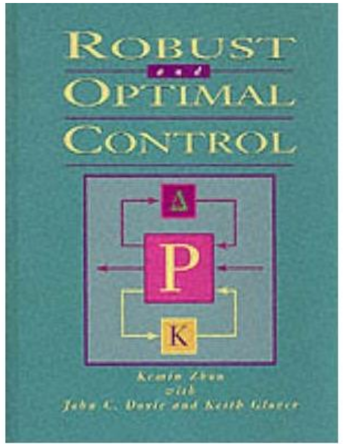

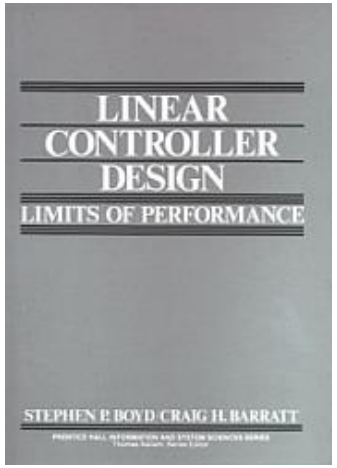

Robust LQG formulation

Nominal LQG formulation

$$\min_{u_0,u_1,\ldots} \quad \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T}\left(y_t^\mathsf{T} Q y_t + u_t^\mathsf{T} R u_t\right)\right]$$
$$\text{subject to} \quad x_{t+1} = A_\star x_t + B_\star u_t + B_\star w_t,$$
$$y_t = C_\star x_t + v_t\,..$$

Robust LQG formulation

$$\min_{\mathbf{K}} \sup_{\|\mathbf{\Delta}\|_\infty<\epsilon} \quad \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T}\left(y_t^\mathsf{T} Q y_t + u_t^\mathsf{T} R u_t\right)\right],$$
$$\text{subject to} \quad \mathbf{y} = (\hat{\mathbf{G}} + \mathbf{\Delta})\mathbf{u} + \mathbf{v}$$
$$\mathbf{u} = \mathbf{K}\mathbf{y} + \mathbf{w},$$

Key idea via Change of variables:

Instead of optimizing the controller K,
we search over the closed-loop responses

Closed-loop convexity

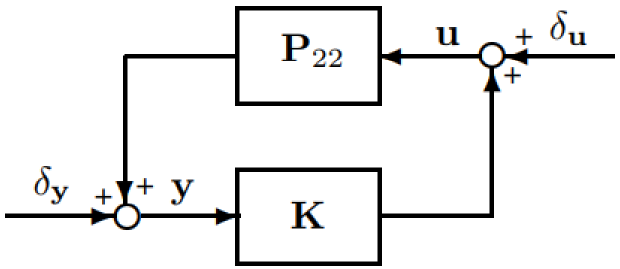

$$\begin{bmatrix}\mathbf{y}\\ \mathbf{u}\end{bmatrix} = \begin{bmatrix}\mathbf{Y} & \mathbf{W}\\ \mathbf{U} & \mathbf{Z}\end{bmatrix}\begin{bmatrix}\delta_y\\ \delta_u\end{bmatrix}$$

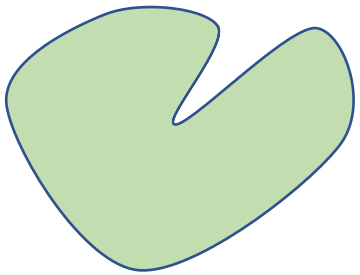

Non-convex

$$\mathbf{K} = \mathbf{U}\mathbf{Y}^{-1}$$

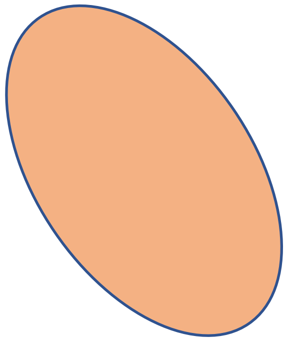

Convex

$$(\mathbf{Y},\mathbf{U},\mathbf{W},\mathbf{Z}) \in \hat{\mathcal{C}}$$
$$\hat{\mathcal{C}} \equiv \text{Affine space} \cap \text{Stable}$$

Furieri, L., Zheng, Y., Papachristodoulou, A., & Kamgarpour, M. (2019). An input–output parametrization of stabilizing controllers: Amidst youla and system level synthesis. *IEEE Control Systems Letters*, *3*(4), 1014-1019.

Robust LQG formulation

Robust LQG formulation

$$\min_{\mathbf{K}} \sup_{\|\boldsymbol{\Delta}\|_\infty < \epsilon} \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T}\left(y_t^\intercal Q y_t + u_t^\intercal R u_t\right)\right],$$

$$\text{subject to} \quad \mathbf{y} = (\hat{\mathbf{G}} + \boldsymbol{\Delta})\mathbf{u} + \mathbf{v}$$

$$\mathbf{u} = \mathbf{K}\mathbf{y} + \mathbf{w},$$

Theorem (Zheng et al., 2021): the problem above is equivalent to

$$\min_{\hat{\mathbf{Y}},\hat{\mathbf{W}},\hat{\mathbf{U}},\hat{\mathbf{Z}}} \sup_{\|\boldsymbol{\Delta}\|_\infty < \epsilon} J(\mathbf{G}_\star, \mathbf{K}) = \left\| \begin{bmatrix} \hat{\mathbf{Y}}(I - \boldsymbol{\Delta}\hat{\mathbf{U}})^{-1} & \hat{\mathbf{Y}}(I - \boldsymbol{\Delta}\hat{\mathbf{U}})^{-1}(\hat{\mathbf{G}} + \boldsymbol{\Delta}) \\ \hat{\mathbf{U}}(I - \boldsymbol{\Delta}\hat{\mathbf{U}})^{-1} & (I - \hat{\mathbf{U}}\boldsymbol{\Delta})^{-1}\hat{\mathbf{Z}} \end{bmatrix} \right\|_{\mathcal{H}_2}$$

$$\text{subject to} \quad \begin{bmatrix} I & -\hat{\mathbf{G}} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{Y}} & \hat{\mathbf{W}} \\ \hat{\mathbf{U}} & \hat{\mathbf{Z}} \end{bmatrix} = \begin{bmatrix} I & 0 \end{bmatrix},$$

$$\begin{bmatrix} \hat{\mathbf{Y}} & \hat{\mathbf{W}} \\ \hat{\mathbf{U}} & \hat{\mathbf{Z}} \end{bmatrix} \begin{bmatrix} -\hat{\mathbf{G}} \\ I \end{bmatrix} = \begin{bmatrix} 0 \\ I \end{bmatrix},$$

$$\hat{\mathbf{Y}}, \hat{\mathbf{W}}, \hat{\mathbf{U}}, \hat{\mathbf{Z}} \in \mathcal{RH}_\infty, \ \|\hat{\mathbf{U}}\|_\infty \leq \frac{1}{\epsilon},$$

Another upper approximation via Taylor expansion

→ Convex optimization

Suboptimality guarantee

Theorem (Zheng et al., 2021): When the plant is open-loop stable, solving an SDP upper approximation of the robust control problem leads to a robust stabilizing LQG control with a suboptimality gap

$$\frac{J(\hat{\mathbf{K}}) - J_\star}{J_\star} \leq 20\epsilon \|\mathbf{U}_\star\|_\infty + \mathcal{O}(\epsilon)\,,$$

where $\|\mathbf{G}_\star - \hat{\mathbf{G}}\|_\infty < \epsilon$, and the estimation is accurate enough

Optimality vs. Robustness

- Certainty equivalent controller (Mania et al., 2019) achieves a better sub-optimality scaling $\mathcal{O}(\epsilon^2)$
- But this method has a much stricter requirement on admissible uncertainty, and has no guarantee of robust stabilization performance

"The price of obtaining a faster rate for LQR is that the certainty equivalent controller becomes less robust to model uncertainty"

- The upper bound depends on the original plant model. Very interesting to see whether certain plants are intrinsically hard to control?

End-to-end Sample complexity

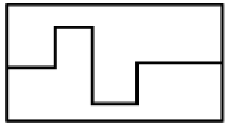

Control inputs u_t

Physical systems

Output meas. y_t

Estimate a model and uncertainty level

$\hat{\mathbf{G}}, \quad \epsilon$

- **Stable system → first T finite impulse responses** (Oymak and Ozay, 2019)

$$\mathbf{G}_\star(z) = \sum_{i=0}^{\infty} \frac{1}{z^i} G_{\star,i} = \sum_{i=0}^{T-1} \frac{1}{z^i} G_{\star,i} + \sum_{i=T}^{\infty} \frac{1}{z^i} G_{\star,i},$$

Markov parameters

$$G_\star = \begin{bmatrix} 0 & C_\star B_\star & \cdots & C_\star A_\star^{T-2} B_\star \end{bmatrix} \in \mathbb{R}^{p \times Tm}.$$

Least-square estimator

$$\hat{G} \in \arg\min_{G} \sum_{t=T}^{\bar{N}} \|y_t - G\bar{u}_t\|_2^2.$$

An estimated plant model

$$\hat{\mathbf{G}} := \sum_{k=0}^{T-1} \frac{1}{z^k} \hat{G}_k.$$

End-to-end Sample complexity

- **Hinf estimation error unbound**

$$\|\mathbf{G}_\star - \hat{\mathbf{G}}\|_\infty = \left\| \sum_{t=0}^{T-1} \left(G_\star(k) - \hat{G}_k\right) \frac{1}{z^k} + \sum_{k=T}^{\infty} G_\star(k) \frac{1}{z^k} \right\|_\infty$$
$$\leq \underbrace{\left\| \sum_{t=0}^{T-1} \left(G_\star(k) - \hat{G}_k\right) \frac{1}{z^k} \right\|_\infty}_{\text{FIR estimation error}} + \underbrace{\left\| \sum_{k=T}^{\infty} G_\star(k) \frac{1}{z^k} \right\|_\infty}_{\text{FIR truncation error}},$$

An estimated plant model

$$\hat{\mathbf{G}} := \sum_{k=0}^{T-1} \frac{1}{z^k} \hat{G}_k.$$

Proposition (Zheng et al., 2021): For open-loop stable plant, with high probability, we have

$$\|\mathbf{G}_\star - \hat{\mathbf{G}}\|_\infty \leq \frac{R_w + R_v + R_e}{\sigma_u} \sqrt{\frac{T}{N}} + \Phi(A_\star)\|C_\star\|\|B_\star\| \frac{\rho(A_\star)^T}{1 - \rho(A_\star)}.$$

- N is the number of samples (y_t, u_t)
- R_w, R_v, R_e are some problem dependent constants
- The last term decreases exponentially to zero as the FIT length T increases

End-to-end Sample complexity

Nominal LQG formulation

$$\min_{u_0, u_1, \ldots} \quad \lim_{T \to \infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T}\left(y_t^\mathsf{T} Q y_t + u_t^\mathsf{T} R u_t\right)\right]$$

$$\text{subject to} \quad x_{t+1} = A_\star x_t + B_\star u_t + B_\star w_t,$$

$$y_t = C_\star x_t + v_t \,.\,.$$

SysID + Robust LQG

$$\min_{\mathbf{K}} \sup_{\|\boldsymbol{\Delta}\|_\infty < \epsilon} \quad \lim_{T \to \infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T}\left(y_t^\mathsf{T} Q y_t + u_t^\mathsf{T} R u_t\right)\right],$$

$$\text{subject to} \quad \mathbf{y} = (\hat{\mathbf{G}} + \boldsymbol{\Delta})\mathbf{u} + \mathbf{v}$$

$$\mathbf{u} = \mathbf{K}\mathbf{y} + \mathbf{w},$$

End-to-end Sample complexity:

Suppose the true plant is FIR of order T_0 and let the length $T \geq T_0$. With high probability, the end-to-end sample complexity scales as

$$\frac{J(\hat{\mathbf{K}}) - J_\star}{J_\star} \sim \mathcal{O}\left(\frac{1}{\sqrt{N}}\right),$$

- N is the number of samples (y_t, u_t) in a single trajectory
- **Robust stability**: as long as the Robust LQG has a feasible solution, the closed-loop is guaranteed to be stable:

Comparison with LQR

	LQR	
	$$\min_{\mathbf{K}} \sup_{\|\Delta_A\|,\|\Delta_B\|<\epsilon} \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T}(x_t^\mathsf{T} Q x_t + u_t^\mathsf{T} R u_t)\right]$$ subject to $x_{t+1} = (\hat{A}+\Delta A)x_t + (\hat{B}+\Delta B)u_t + v_t$, $\mathbf{u} = \mathbf{K}\mathbf{x}$	$$\min_{\mathbf{K}} \sup_{\|\mathbf{\Delta}\|_\infty<\epsilon} \lim_{T\to\infty} \mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T}(y_t^\mathsf{T} Q y_t + u_t^\mathsf{T} R u_t)\right]$$ subject to $\mathbf{y} = (\hat{\mathbf{G}}+\mathbf{\Delta})\mathbf{u} + \mathbf{v}$, $\mathbf{u} = \mathbf{K}\mathbf{y} + \mathbf{w}$,
Sys ID methods	❖ Least squares $\|\hat{A} - A_\star\| \le \epsilon_A, \|\hat{B} - B_\star\| \le \epsilon_B,$	❖ Least squares $\|\mathbf{\Delta}\|_\infty := \|\mathbf{G}_\star - \hat{\mathbf{G}}\|_\infty < \epsilon$
Synthesis Technique	❖ Frequency domain ❖ System-level synthesis, SLS (Wang et al., 2019) ❖ Taylor expansion	❖ Frequency domain ❖ Input-output parameterization, IOP, (Furieri et al., 2019) ❖ Taylor expansion
Sample Complexity	❖ both stable and unstable systems $\frac{J(\hat{K}) - J_\star}{J_\star} \sim \mathcal{O}\left(\frac{1}{\sqrt{N}}\right),$	❖ Only for open-loop stable system $\frac{J(\hat{\mathbf{K}}) - J_\star}{J_\star} \sim \mathcal{O}\left(\frac{1}{\sqrt{N}}\right),$
References	✓ Dean et al., 2020; Berberich et al., 2020; Boczar et al., 2018; Tsiamis et al., 2020; Umenberger et al., 2019; and many others	• Zheng, Furieri, Kamgarpour, & Li, (2021, May). link

Conclusion

Two main takeaways

Landscape Analysis of LQG control

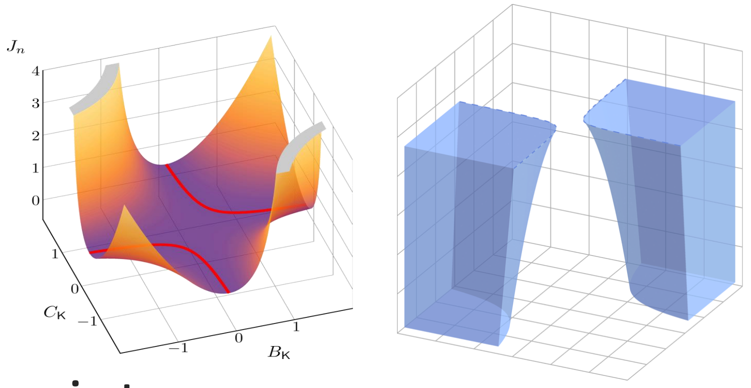

- Much richer and more complicated than LQR
- Disconnected, but at most 2 connected components
- Non-unique, non-isolated stationary points, strict saddle points
- Minimal stationary points are globally optimal

Sample Complexity of LQG control

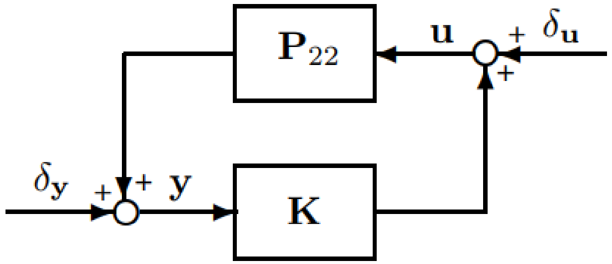

- Robust LQG formulation for stability/safety guarantees
- End-to-end sample complexity is comparable to LQR
- Frequency domain design methods (SLS, IOP) are very useful for learning-based control

Join the SOC lab at UC San Diego!

SOC Lab

Data-driven and learning-based control

Sparse conic optimization

Scalable distributed control

Connected and autonomous vehicles (CAVs)

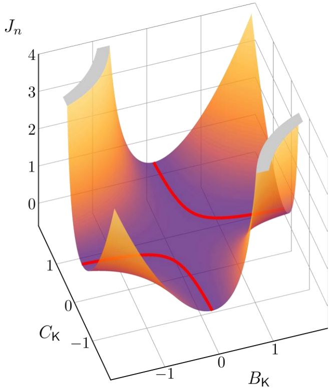

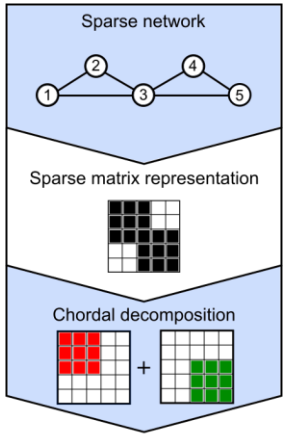

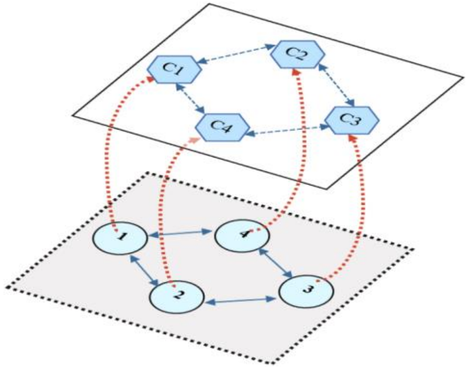

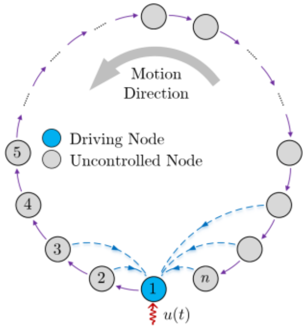

Check out our webpage: https://zhengy09.github.io/soclab.html

Thank you for your attention!

Q & A

More details. Check out our webpage: https://zhengy09.github.io/soclab.html