AFOSR DC Review Meeting 2024

Stability Margins of Neural Network Controllers

Murat Arcak
UC Berkeley
with Neelay Junnarkar and Pete Seiler

Award # FA9550-21-1-0288

Motivation

- Emerging energy efficient airframes pose complex dynamics, not amenable to traditional model reduction and control.
- Low bandwidth control is standard approach to avoid exciting unmodeled dynamics, but it sacrifices performance/agility.
- Neural Network controllers can overcome these difficulties. Favored by several Advanced Air Mobility companies, but…

Challenge: Formal guarantees of stability and sufficient margins

e.g., Code of Federal Regulations on Airworthiness Standards (14 CFR 25) states: "The airplane must be longitudinally, directionally, and laterally stable in accordance with the provisions of §25.173 through 25.177." Advisory Circular on aeroelastic stability demands 6 dB gain, 60° phase margin.

Outline

- Neural Network (NN) Control System

  A unifying recurrent implicit NN architecture for controlling an uncertain plant

- Control Synthesis

  NN training to maximize a reward subject to a dissipativity constraint to accommodate a class of plant uncertainties

- Examples

  1. Inverted pendulum, 2. Flexible rod on a cart

- Disk Margins

  Incorporating classical gain/phase margins and elevating them to nonlinear systems through dissipativity theory

NN Control System

Recurrent Implicit Neural Network (in continuous time):

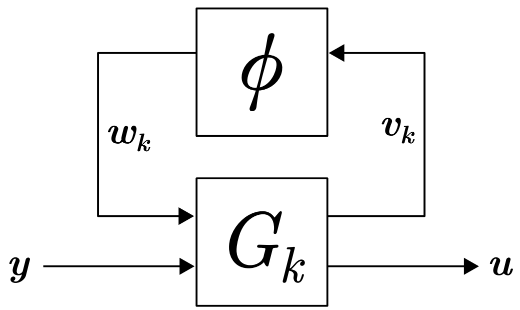

$$\begin{bmatrix} \dot{\boldsymbol{x}}_{\boldsymbol{k}}(t) \\ \boldsymbol{v}_{\boldsymbol{k}}(t) \\ \boldsymbol{u}(t) \end{bmatrix} = \begin{bmatrix} A_k & B_{kw} & B_{ky} \\ C_{kv} & D_{kvw} & D_{kvy} \\ C_{ku} & D_{kuw} & D_{kuy} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}_{\boldsymbol{k}}(t) \\ \boldsymbol{w}_{\boldsymbol{k}}(t) \\ \boldsymbol{y}(t) \end{bmatrix}$$

$$\boldsymbol{w}_{\boldsymbol{k}}(t) = \phi(\boldsymbol{v}_{\boldsymbol{k}}(t)),$$

G_k : LTI model for the linear dynamics of the controller

ϕ : vector of activation functions, acting componentwise

"Implicit" due to equation: $w_k = \phi(C_{kv}x_k + D_{kvw}w_k + D_{kvy}y)$

Can be made explicit in special cases, e.g. feedforward networks where D_{kvw} is strictly upper triangular. Our training procedure ensures well-posedness without restricting the network structure.

NN Control System

Closed-loop with plant G_p and model uncertainty Δ_p:

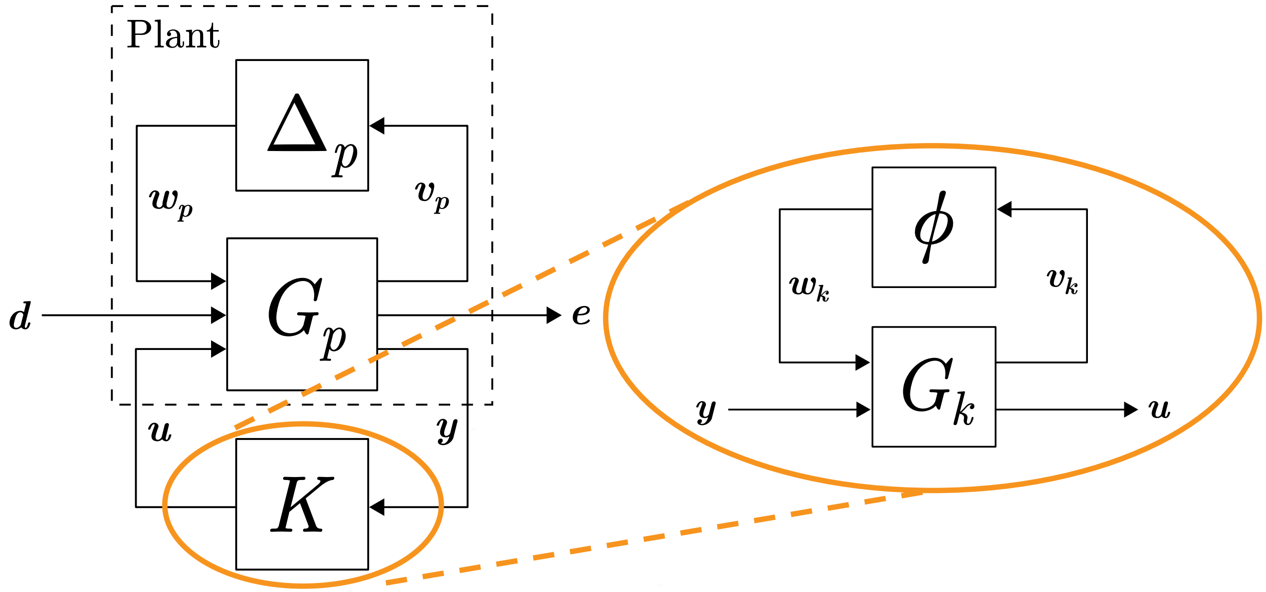

In standard LFT form:

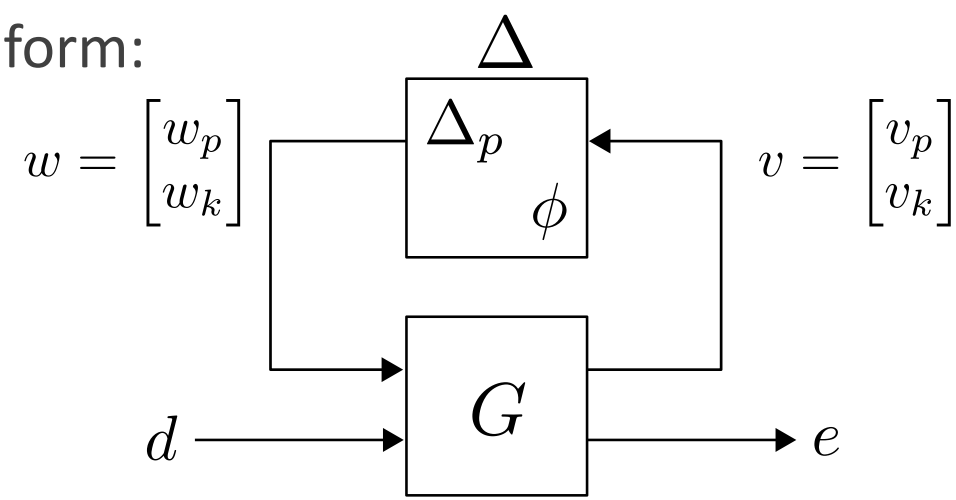

NN Control System

1. Characterize activation functions and plant uncertainty with Integral Quadratic Constraints (IQCs)
2. Formulate closed-loop stability requirement as dissipation inequality to be satisfied by nominal system G
3. Train NN controller (i.e. weights in G_k) to maximize a reward subject to the dissipativity constraint

Static IQCs:

$$\int_0^T \begin{bmatrix} v(t) \\ w(t) \end{bmatrix}^\top M \begin{bmatrix} v(t) \\ w(t) \end{bmatrix} dt \geq 0 \quad \forall T \geq 0$$

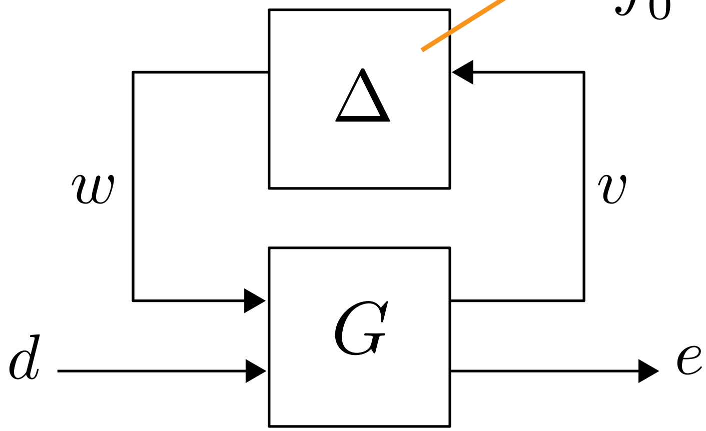

Simplest quadratic constraint for activation functions in sector [0,1]:

$$M_\phi = \begin{bmatrix} 0 & \Lambda \\ \Lambda & -2\Lambda \end{bmatrix}, \quad \Lambda \succeq 0 \text{ diagonal}$$

NN Control System

Dynamic IQCs: $$\int_0^T \begin{bmatrix} \tilde{v}(t) \\ \tilde{w}(t) \end{bmatrix}^\top M \begin{bmatrix} \tilde{v}(t) \\ \tilde{w}(t) \end{bmatrix} dt \geq 0 \quad \forall T \geq 0$$

where $\tilde{v}, \tilde{w}$ are outputs of "filters" Ψ_1, Ψ_2 applied to v, w.

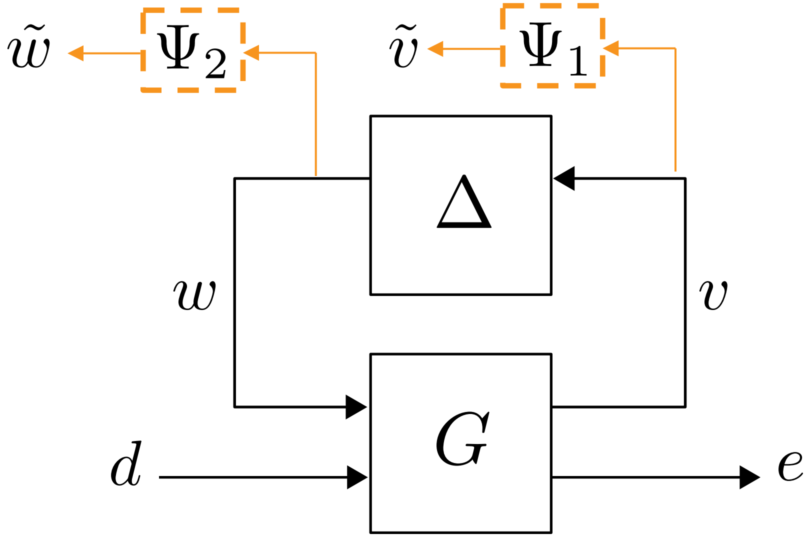

More expressive than static IQCs: can choose Ψ_1, Ψ_2, e.g., to emphasize different frequency ranges in signals v, w.

NN Control System

2. Formulate closed-loop stability requirement as dissipation inequality to be satisfied by nominal system G

Dissipativity: Take combined dynamical model for G, Ψ_1, Ψ_2:

$$\dot{x}(t) = f(x(t), w(t), d(t))$$

If there exists nonnegative storage function $x \mapsto V(x)$ such that

$$\nabla V(x)^\top f(x, w, d) + \begin{bmatrix} \tilde{v} \\ \tilde{w} \end{bmatrix}^\top M \begin{bmatrix} \tilde{v} \\ \tilde{w} \end{bmatrix} \leq s(d, e) \quad \forall x, w, d$$

then $V(x(T)) \leq V(x(0)) + \int_0^T s(d(t), e(t)) dt \quad \forall T \geq 0$

This guarantees Lyapunov stability with $s = 0$, L_2 gain with $s(d, e) = \gamma^2 |d|^2 - |e|^2$, etc., for any Δ satisfying the IQC.

Control Synthesis

3. Train NN controller (i.e. weights in G_k) to maximize a reward subject to the dissipativity constraint

- Training algorithm alternates between a Reinforcement Learning (RL) step and a dissipativity-enforcing step.
- RL step aims to maximize reward $\int_0^T r(x(t),u(t))dt$ averaged over disturbances and initial conditions seen during training.
- Dissipativity-enforcing step computationally tractable for linear G (i.e., if plant nonlinearities can be subsumed in Δ):

$$G, \Psi_1, \Psi_2 : \quad \begin{bmatrix} \dot{\boldsymbol{x}}(t) \\ \boldsymbol{v}(t) \\ \boldsymbol{e}(t) \end{bmatrix} = \begin{bmatrix} A & B_w & B_d \\ C_v & D_{vw} & D_{vd} \\ C_e & D_{ew} & D_{ed} \end{bmatrix} \begin{bmatrix} \boldsymbol{x}(t) \\ \boldsymbol{w}(t) \\ \boldsymbol{d}(t) \end{bmatrix}$$

Control Synthesis

- Matrices above are affine in control parameters:

$$\theta \triangleq \begin{bmatrix} A_k & B_{kw} & B_{ky} \\ C_{kv} & D_{kvw} & D_{kvy} \\ C_{ku} & D_{kuw} & D_{kuy} \end{bmatrix}.$$

- With quadratic storage function $V(x) = x^\top P x$, the dissipativity condition becomes a Bilinear Matrix Inequality in P, θ, Λ where Λ is the diagonal matrix in M_ϕ.
- Change of variables to obtain Linear Matrix Inequality (LMI) in new decision variables, from which P, θ, Λ recovered.
- LMI incorporates $\Lambda D_{kvw} + D_{kvw}^\top \Lambda - 2\Lambda \preceq 0.$ Imposing strictness of this inequality ensures well-posedness of the implicit NN for activation functions with slope bound [0,1].

Examples

1. Stabilizing an inverted pendulum with minimal control effort

$$\begin{aligned}\dot{\boldsymbol{x}}_{\mathbf{1}}(t) &= \boldsymbol{x}_{\mathbf{2}}(t)\\ \dot{\boldsymbol{x}}_{\mathbf{2}}(t) &= -\frac{\mu}{m\ell^2}\boldsymbol{x}_{\mathbf{2}}(t) + \frac{g}{\ell}\underbrace{\sin(\boldsymbol{x}_{\mathbf{1}}(t))}_{\Delta_p,\ \text{covered w/ sector bound}} + \frac{1}{m\ell^2}\boldsymbol{u}(t)\\ \boldsymbol{y}(t) &= \boldsymbol{x}_{\mathbf{1}}(t)\end{aligned}$$

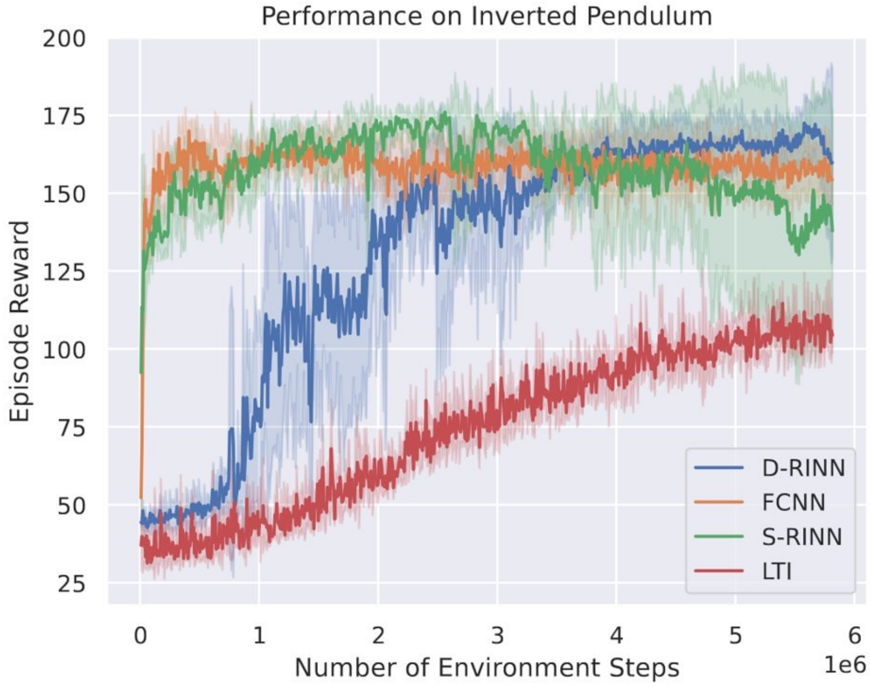

D-RINN: RINN with dissipativity constraint

S-RINN: RINN w/out dissipativity constraint

FCNN: Fully connected NN

LTI controller w/ dissipativity constraint, trained like D-RINN

Examples

2. Flexible rod on a cart

Nominal model for rigid rod + uncertainty with L_2 gain bound cover flexible model.

Training with flexible model to minimize cost in state and control.

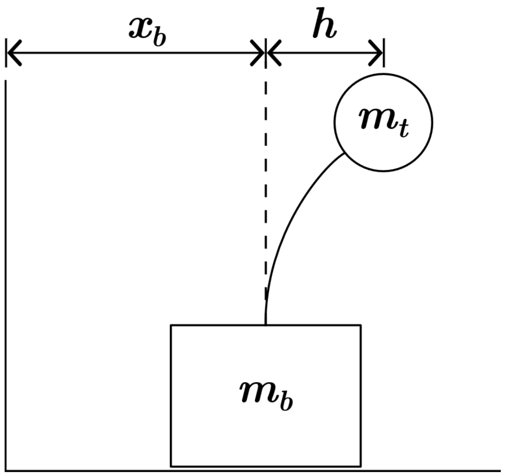

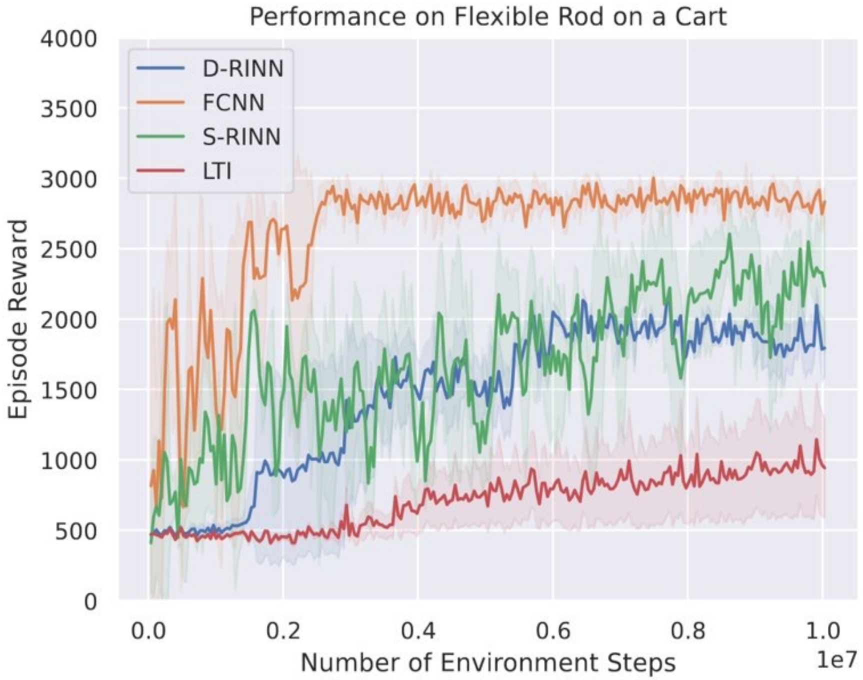

D-RINN: RINN with dissipativity constraint

S-RINN: RINN w/out dissipativity constraint

FCNN: Fully connected NN

LTI controller w/ dissipativity constraint, trained like D-RINN

Disk Margins

Classical definition: Linear SISO plant G_p with controller K has disk margin $D(\alpha,\sigma)$ if $1-\delta\mathbf{G_p}(\omega)\mathbf{K}(\omega)\neq 0$

$$\forall\omega,\ \forall\delta\in D(\alpha,\sigma):=\left\{\frac{1+\frac{1-\sigma}{2}z}{1-\frac{1+\sigma}{2}z}:z\in\mathbb{C},|z|\leq\alpha\right\}$$

Generalization: G_p with controller K has disk margin $D(\alpha,\sigma)$ if perturbed system below is stable for any Δ_p with L_2 gain $\leq\alpha$:

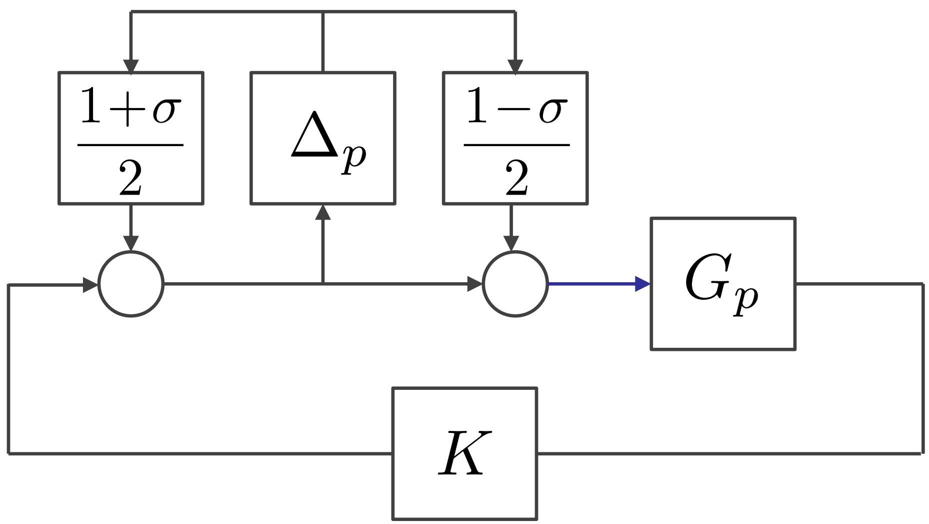

Disk Margins

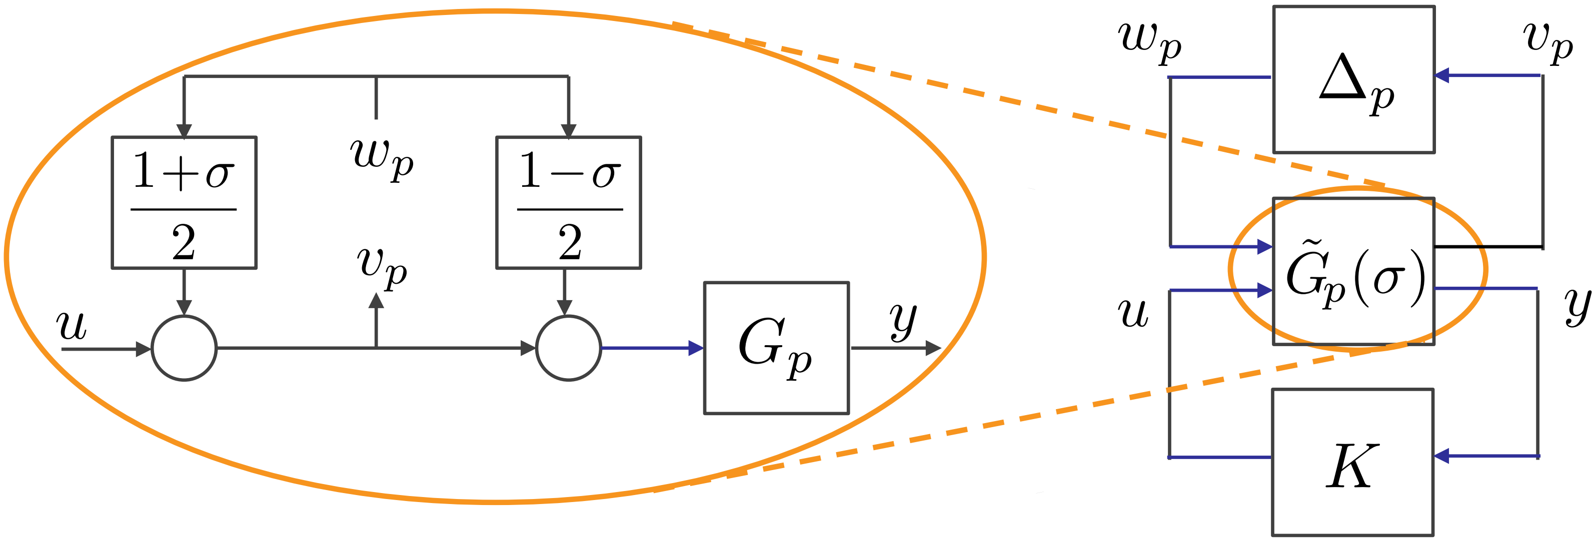

Final LFT, merging LTI part of NN controller with $\tilde{G}_p(\sigma)$ into $\tilde{G}(\sigma)$:

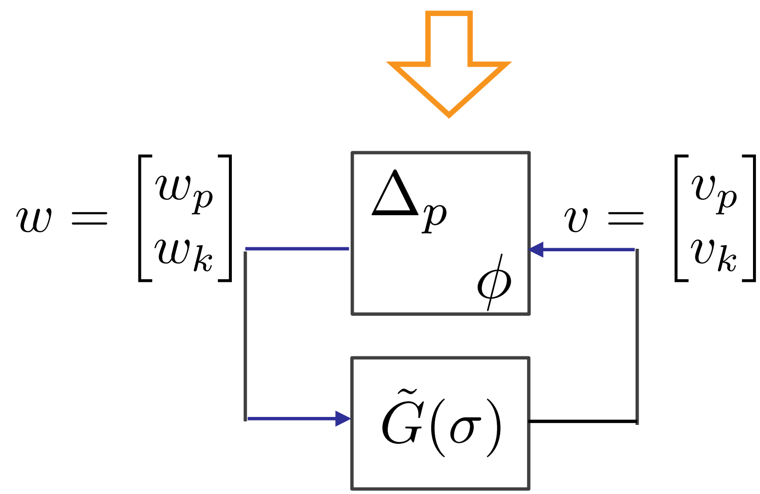

Δ_p satisfies IQC for L_2 gain $\leq \alpha$:

$$M_p = \begin{bmatrix} \alpha^2 I & 0 \\ 0 & -I \end{bmatrix}$$

Synthesis method applied to this setup guarantees disk margin.

Related Publications

1. Yin, Seiler, Arcak. Stability analysis using quadratic constraints for systems with neural network controllers. IEEE TAC, 2022.
2. Yin, Seiler, Jin, Arcak. Imitation learning with stability and safety guarantees. IEEE LCSS, 2022.
3. Gu, Yin, El Ghaoui, Arcak, Seiler, Jin. Recurrent neural network controller synthesis with stability guarantees for partially observed systems. AAAI Conf. on Artificial Intelligence, 2022.
4. Junnarkar, Yin, Gu, Arcak, Seiler. Synthesis of stabilizing recurrent equilibrium network controllers. IEEE CDC 2022.
5. Junnarkar, Arcak, Seiler. Synthesizing neural network controllers with closed-loop dissipativity guarantees. arXiv:2404.07373v1
6. Junnarkar, Arcak, Seiler. Stability margins of neural network controllers. Submitted.