Efficient Inference of Interaction Graphs in Nonlinear Multi-Agent Networks

IFAC Pre-Congress Workshop: Alternative Representations for Control

Miel Sharf and **Daniel Zelazo**

The Stephen B. Klein Faculty of Aerospace Engineering, Technion

August 23, 2026

CONNECT LAB
Cooperative Networks and Controls

THE STEPHEN B. KLEIN FACULTY OF AEROSPACE ENGINEERING

TECHNION Israel Institute of Technology

Networks Are Part of the Model

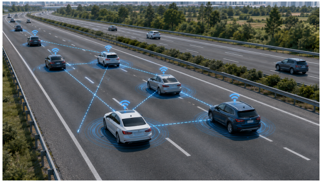

Distributed autonomy

Biological networks

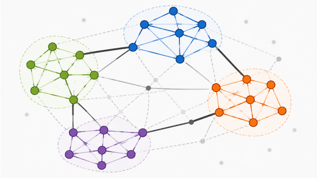

Network science

- **Distributed autonomy:** communication networks coordinate autonomous agents.
- **Biological networks:** interaction networks encode influence among biological components.
- **Network science:** weighted graphs organize and shape collective behavior.

> *Across very different domains, the same object appears:*
> *the **interaction graph** governing who influences whom.*

Networks Are Part of the Model

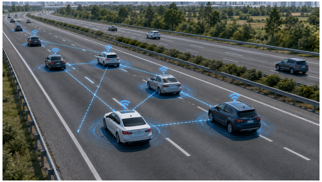

Distributed autonomy

Biological networks

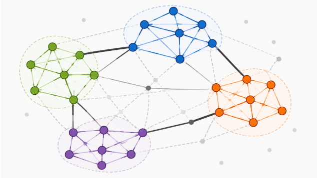

Network science

- **Distributed autonomy:** communication networks coordinate autonomous agents.
- **Biological networks:** interaction networks encode influence among biological components.
- **Network science:** weighted graphs organize and shape collective behavior.

Question

Can we efficiently infer the interaction graph of a **nonlinear network** from controlled experiments?

What Should We Identify?

Workshop viewpoint: the right representation reveals what can be inferred from data.

What Should We Identify?

Workshop viewpoint: the right representation reveals what can be inferred from data.

Transient-dynamics view

Identify the evolution

$$\dot{x} = f(x, u; \mathcal{G})$$

Use trajectories, correlations,
time-series models, or local linearizations:

$$\{x(t), u(t)\}_{t \geq 0} \quad \rightarrow \quad \mathcal{G}.$$

Difficulty: nonlinear, nonsmooth,
or discontinuous dynamics.

Materassi et al., TAC 2010; Mauroy et al., CDC 2017;

Li et al., Chaos 2015; Burbano Lombana et al., TCNS 2020.

What Should We Identify?

Workshop viewpoint: the right representation reveals what can be inferred from data.

Transient-dynamics view

Identify the evolution

$$\dot{x} = f(x, u; \mathcal{G})$$

Use trajectories, correlations, time-series models, or local linearizations:

$$\{x(t), u(t)\}_{t \geq 0} \quad \to \quad \mathcal{G}.$$

Difficulty: nonlinear, nonsmooth, or discontinuous dynamics.

Steady-state response view

Identify the equilibrium map

$$w \mapsto y_{\mathrm{ss}}(w)$$

Apply constant probing inputs and observe the converged outputs:

$$\{w^{(k)}, y_{\mathrm{ss}}^{(k)}\}_{k=0}^{n} \quad \to \quad \mathcal{G}.$$

Opportunity: the graph appears in the sensitivity of this map.

What Should We Identify?

Workshop viewpoint: the right representation reveals what can be inferred from data.

Transient-dynamics view

Identify the evolution

$$\dot{x} = f(x, u; \mathcal{G})$$

Use trajectories, correlations, time-series models, or local linearizations:

$$\{x(t), u(t)\}_{t \geq 0} \quad \to \quad \mathcal{G}.$$

Difficulty: nonlinear, nonsmooth, or discontinuous dynamics.

Steady-state response view

Identify the equilibrium map

$$w \mapsto y_{\text{ss}}(w)$$

Apply constant probing inputs and observe the converged outputs:

$$\{w^{(k)}, y_{\text{ss}}^{(k)}\}_{k=0}^{n} \quad \to \quad \mathcal{G}.$$

Opportunity: the graph appears in the sensitivity of this map.

Guiding idea: *Replace nonlinear transient-dynamics identification by structured inference from the* ***steady-state input-output representation***.

Diffusively Coupled Nonlinear Networks

Closed-loop architecture

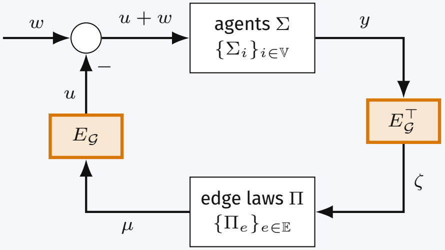

Local models

Each vertex $i \in \mathbb{V}$ is a dynamical agent

$$\Sigma_i : \quad \begin{cases} \dot{x}_i = f_i(x_i) + q_i(x_i)(u_i + w_i), \\ y_i = h_i(x_i). \end{cases}$$

Single-valued monotone edge law

$$\Pi_e : \quad \mu_e = g_e(\zeta_e).$$

The agent and interaction models are known; the way they are interconnected is not.

Diffusively Coupled Nonlinear Networks

Closed-loop architecture

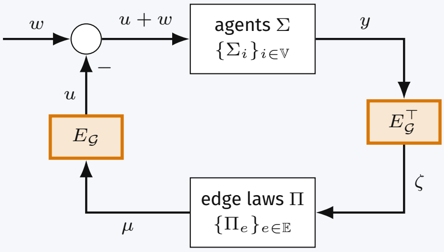

The incidence matrix, $E_{\mathcal{G}}$, routes **relative outputs** to edge laws and returns their actions to the agents.

Diffusive interconnection

Interaction graph

$$\mathcal{G} = (\mathbb{V}, \mathbb{E}), \qquad E_{\mathcal{G}} \in \mathbb{R}^{|\mathbb{V}| \times |\mathbb{E}|} \quad \text{incidence matrix}$$

Relative outputs

$$\zeta = E_{\mathcal{G}}^{\top} y$$

Weighted edge interactions

$$\mu = N g(\zeta), \qquad N = \mathrm{diag}(\nu_e)$$

Diffusive feedback

$$u = -E_{\mathcal{G}} \mu$$

The unknown graph enters the dynamics entirely through $E_{\mathcal{G}}$ and the edge weights ν_e.

Diffusively Coupled Nonlinear Networks

Closed-loop architecture

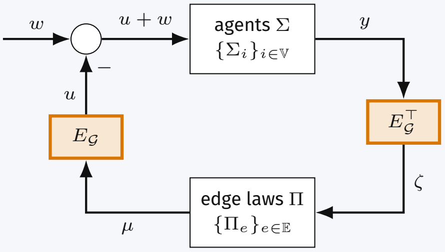

The incidence matrix, $E_{\mathcal{G}}$, routes **relative outputs** to edge laws and returns their actions to the agents.

Diffusive interconnection

Interaction graph

$$\mathcal{G} = (\mathbb{V},\mathbb{E}), \qquad E_{\mathcal{G}} \in \mathbb{R}^{|\mathbb{V}|\times|\mathbb{E}|} \quad \text{incidence matrix}$$

Relative outputs

$$\zeta = E_{\mathcal{G}}^\top y$$

Weighted edge interactions

$$\mu = N g(\zeta), \qquad N = \mathrm{diag}(\nu_e)$$

Diffusive feedback

$$u = -E_{\mathcal{G}}\mu$$

The unknown graph enters the dynamics entirely through $E_{\mathcal{G}}$ and the edge weights ν_e.

Structure learning with known local physics: *we know how the components behave, but not **which components are actually coupled**.*

Identification Strategy: Probe, Settle, Infer

Use constant inputs to turn the dynamical network into a collection of steady-state experiments.

choose probes $\{\bar{w}^{(k)}\}_{k=0}^{m}$ → **apply each probe** to the nonlinear network → **wait and record** $\mathcal{D} = \{(\bar{w}^{(k)}, \bar{y}^{(k)})\}$ → **use equilibrium shifts** $(\delta w^{(k)}, \delta y^{(k)}) \Rightarrow \widehat{\mathcal{G}}$

What the experiment controls

The constant probe values $\bar{w}^{(k)}$, chosen near a baseline operating point.

What identification uses

Changes between the resulting equilibria—not a model of the full transient trajectory.

Identification Strategy: Probe, Settle, Infer

Use constant inputs to turn the dynamical network into a collection of steady-state experiments.

choose probes $\{\bar{w}^{(k)}\}_{k=0}^{m}$ → **apply each probe** to the nonlinear network → **wait and record** $\mathcal{D} = \{(\bar{w}^{(k)}, \bar{y}^{(k)})\}$ → **use equilibrium shifts** $(\delta w^{(k)}, \delta y^{(k)}) \Rightarrow \widehat{\mathcal{G}}$

What the experiment controls

The constant probe values $\bar{w}^{(k)}$, chosen near a baseline operating point.

What identification uses

Changes between the resulting equilibria—not a model of the full transient trajectory.

The entire strategy rests on each probe producing a trustworthy steady-state pair.

Before Identification: Can We Trust Each Steady-State Measurement?

A pair $(\bar{w}, \bar{y})$ is useful only when three separate requirements are satisfied.

Existence

The steady-state equations admit a solution for $\bar{w}$.

certificate: feasibility

$+$

Convergence

The trajectory settles to an equilibrium.

certificate: passivity / MEIP

$+$

Usable Data

A finite waiting time guarantees the desired accuracy.

certificate: profile bound

Together they certify one valid experiment

$\bar{w} \mapsto y(t) \to \bar{y} \quad \Rightarrow \quad (\bar{w}, \bar{y})$ can be used for identification.

We now establish these certificates, beginning with network convergence.

Passivity About a Fixed Equilibrium Certifies Convergence

Fix a known equilibrium $(\bar{u}, \bar{y}, \bar{\zeta}, \bar{\mu})$ generated by a constant input $\bar{w}$.

Passivity relative to this equilibrium

Write $\widetilde{y} = y - \bar{y}$, and similarly for the other signals. Assume each agent is output-strictly passive:

$$\dot{S}_i \le \widetilde{u}_i \widetilde{y}_i - \rho_i |\widetilde{y}_i|^2, \qquad \rho_i > 0,$$

and the weighted edge map is monotone:

$$\widetilde{\mu}^\top \widetilde{\zeta} \ge 0.$$

Proposition (convergence to $\bar{y}$) [Arcak TAC2007]

The diffusive interconnection gives

$$\widetilde{y}^\top \widetilde{u} = -\widetilde{\zeta}^\top \widetilde{\mu} \le 0.$$

Hence the summed storage satisfies

$$\dot{S}_{\text{net}} \le -\sum_i \rho_i |\widetilde{y}_i|^2.$$

If trajectories are bounded and the standard invariance conditions hold, then $y(t) \to \bar{y}$.

The Nuance: A Network Has Many Equilibria

The equilibrium reached by the network depends on the constant probe that is applied.

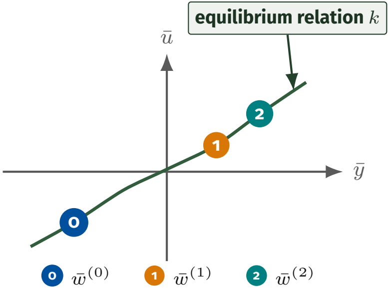

Passivity at one operating point

A standard certificate is written relative to one chosen equilibrium:

$$\dot{S} \leq (u-\bar{u})(y-\bar{y}) - \rho|y-\bar{y}|^2.$$

But probing moves the target

$$\bar{w}^{(0)} \mapsto \bar{y}^{(0)}, \qquad \bar{w}^{(1)} \mapsto \bar{y}^{(1)}, \qquad \bar{w}^{(2)} \mapsto \bar{y}^{(2)}.$$

The relevant equilibrium is probe-dependent and generally not known before the experiment.

The Nuance: A Network Has Many Equilibria

The equilibrium reached by the network depends on the constant probe that is applied.

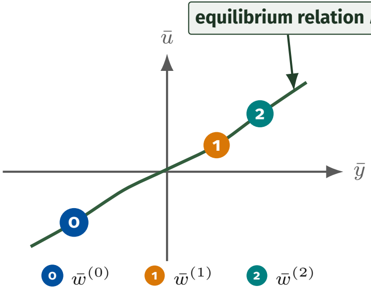

Passivity at one operating point

A standard certificate is written relative to one chosen equilibrium:

$$\dot{S} \leq (u - \bar{u})(y - \bar{y}) - \rho|y - \bar{y}|^2.$$

But probing moves the target

$$\bar{w}^{(0)} \mapsto \bar{y}^{(0)}, \qquad \bar{w}^{(1)} \mapsto \bar{y}^{(1)}, \qquad \bar{w}^{(2)} \mapsto \bar{y}^{(2)}.$$

The relevant equilibrium is probe-dependent and generally not known before the experiment.

We need a convergence certificate valid across the entire equilibrium family.

EIP and MEIP: Passivity Across the Equilibrium Set

This is the refinement needed when a probing experiment changes the operating point.

EIP: every equilibrium **[Hines et al. AUT2011]**

The steady-state relation k_i collects the equilibrium pairs

$$(\bar{u}_i, \bar{y}_i) \in k_i.$$

For **every** pair in k_i, there is a storage function satisfying

$$\dot{S}_i \leq (u_i - \bar{u}_i)(y_i - \bar{y}_i) - \rho_i |y_i - \bar{y}_i|^2.$$

MEIP: the maximal relation

[Bürger, Z, Allogwer AUT2014]

MEIP adds that k_i is a **maximally monotone** relation.

Monotone: equilibrium input-output pairs preserve their ordering.

Maximal: no additional pair can be included without losing monotonicity.

Network conclusion: *feasibility + MEIP $\Rightarrow y(t) \to \bar{y}$ for each constant probe.*

The Network as a Steady-State Map

At equilibrium, the nonlinear network is described by
a static relation between constant inputs and outputs.

Equilibrium equations

Agent steady-state relation:

$$k^{-1}(\bar{y}) = \bar{u} + \bar{w}.$$

Diffusive interconnection:

$$\bar{\zeta} = E_{\mathcal{G}}^{\top}\bar{y}, \qquad \bar{\mu} = Ng(\bar{\zeta}),$$

$$\bar{u} = -E_{\mathcal{G}}\bar{\mu}.$$

Eliminate internal variables

Substitute the interconnection into the steady-state relation:

$$k^{-1}(\bar{y}) = -E_{\mathcal{G}}Ng(E_{\mathcal{G}}^{\top}\bar{y}) + \bar{w}.$$

Equivalently,

$$\boxed{\bar{w} = k^{-1}(\bar{y}) + E_{\mathcal{G}}Ng(E_{\mathcal{G}}^{\top}\bar{y})}$$

Alternative representation: *infer $\mathcal{G}$ from the steady-state map $\bar{w} \mapsto \bar{y}$, rather than from transient trajectories.*

The Network as a Steady-State Map

At equilibrium, the nonlinear network is described by
a static relation between constant inputs and outputs.

Equilibrium equations

Agent steady-state relation:

$$k^{-1}(\bar{y}) = \bar{u} + \bar{w}.$$

Diffusive interconnection:

$$\bar{\zeta} = E_{\mathcal{G}}^{\top}\bar{y}, \qquad \bar{\mu} = Ng(\bar{\zeta}),$$

$$\bar{u} = -E_{\mathcal{G}}\bar{\mu}.$$

Response-map viewpoint

Define the graph-dependent map

$$F_{\mathcal{G}}(\bar{y}) := k^{-1}(\bar{y}) + E_{\mathcal{G}}Ng(E_{\mathcal{G}}^{\top}\bar{y}).$$

Then a steady-state experiment gives

$$\bar{w} = F_{\mathcal{G}}(\bar{y}), \qquad \bar{w} \mapsto \bar{y}.$$

The unknown graph is encoded in
the structure of $F_{\mathcal{G}}$.

Alternative representation: *infer $\mathcal{G}$ from the steady-state map $\bar{w} \mapsto \bar{y}$, rather than from transient trajectories.*

A Convergence Profile Gives a Finite Waiting Time

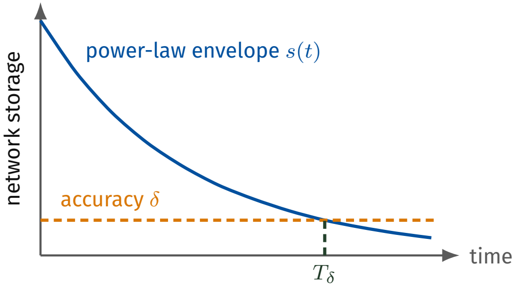

The curve is a conservative bound,

not a model of the full transient.

Power-law comparison envelope

For suitable $C > 0$ and $0 < \beta < 2$, the network storage can be bounded by

$$S_{\text{net}}(t) \leq s(t), \qquad \dot{s} = -Cs^{\beta},$$
$$s(0) = S_{\text{net}}(0).$$

[Sharf, Z,TAC 2022]

Implication for an experiment

Choose δ to imply the desired output accuracy, and define

$$T_\delta := \inf\{t \geq 0 : s(t) \leq \delta\}.$$

Then $T \geq T_\delta$ is a certified hold time.

The Certificates Make Each Probe Trustworthy

✓ Existence

The steady-state equation

$$\bar{w} = F_{\mathcal{G}}(\bar{y})$$

admits an equilibrium output.

✓ Convergence

Passivity across the equilibrium family ensures

$$y(t) \longrightarrow \bar{y}.$$

The recorded limit is meaningful.

✓ Finite waiting

The convergence profile gives a hold time:

$$T \geq T_{\delta}.$$

The requested accuracy is certified.

choose $\bar{w}$ → apply to network → wait T_{δ} → record $\bar{y}$

We can now treat the nonlinear network as a reliable steady-state oracle $\bar{w} \mapsto \bar{y}$ and ask how nearby probes reveal its graph.

How Do Probes Reveal the Graph?

1. Record a baseline

Apply a constant input, wait, and record the equilibrium output:

$$w^{(0)} \xrightarrow{\text{wait}} \bar{y}^{(0)}.$$

2. Probe node i

In experiment i,

$$w^{(i)} = w^{(0)} + \kappa e_i.$$

$[e_i]_k = 0$ if $k \neq i$, and $[e_i]_k = 1$ if $k = i$.
After waiting, record $\delta y^{(i)} := \bar{y}^{(i)} - \bar{y}^{(0)}$.

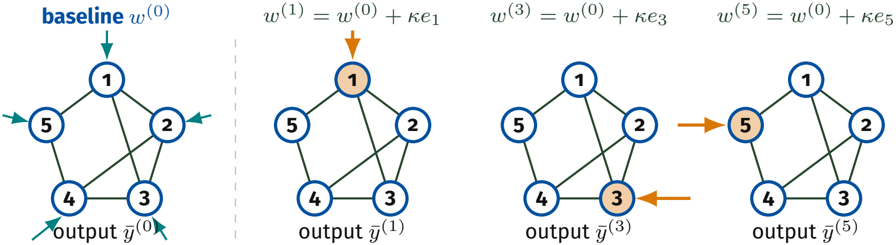

teal arrows: components of $w^{(0)}$ orange: added probe $+\kappa e_i$ blue outline: measured output

From Probe Responses to a Local Model

At equilibrium, $w = F_{\mathcal{G}}(y)$; write $\bar{y} := \bar{y}^{(0)}$ and estimate $\mathcal{M} := DF_{\mathcal{G}}(\bar{y})$.

One experiment: expose the Taylor remainder

Subtract baseline and expand around $\bar{y}$:

$$\delta w^{(i)} = \mathcal{M}\delta y^{(i)} + r^{(i)}.$$

For a twice differentiable equilibrium map,

$$\|r^{(i)}\| = O\left(\|\delta y^{(i)}\|^2\right).$$

The remainder is caused by curvature, even with exact settling and noiseless sensors.

Stack the experiments and estimate the Jacobian

With $\delta W = \kappa I$, define

$$\delta Y = [\,\delta y^{(1)} \ \cdots \ \delta y^{(n)}\,], \quad R = [\,r^{(1)} \ \cdots \ r^{(n)}\,].$$

Then $\boxed{\delta W = \mathcal{M}\delta Y + R}$.

If δY is nonsingular,

$$\widehat{\mathcal{M}} := \delta W \delta Y^{-1}, \qquad \boxed{\widehat{\mathcal{M}} - \mathcal{M} = R\delta Y^{-1}}.$$

Why Does $\mathcal{M}$ Reveal the Graph?

$$\mathcal{M} = \underbrace{\nabla k^{-1}(\bar{y})}_{\text{diagonal agent term}} + \underbrace{E_{\mathcal{G}} D E_{\mathcal{G}}^{\top}}_{\text{weighted graph Laplacian}}, \qquad D = \operatorname{diag}(\nu_e d_e).$$

$d_{ij} := g'_{ij}(\bar{y}_i - \bar{y}_j)$ is the local edge-law slope.

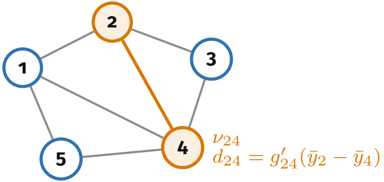

network $\mathcal{G}$

edge $\{2, 4\}$
$\mathcal{M}_{24} = \mathcal{M}_{42}$
$= -\nu_{24} d_{24}$

Jacobian $\mathcal{M}$

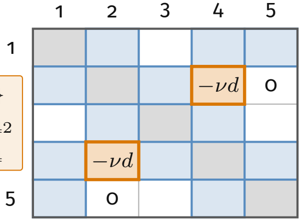

gray: agent terms blue/orange: edges white: non-edges

For $i \neq j$,

$$\boxed{\mathcal{M}_{ij} = \begin{cases} -\nu_{ij} d_{ij}, & \{i, j\} \in \mathbb{E}, \\ 0, & \{i, j\} \notin \mathbb{E}. \end{cases}}$$

Hence topology is recovered when $\nu_{ij} d_{ij} \neq 0$.

Why Do Small Probes Produce a Basis?

independent input probes

$$\delta W = \kappa[\, e_1 \ \cdots \ e_n \,] = \kappa I$$

local inverse map

$$DF_{\mathcal{G}}^{-1}(\bar{w}) = \mathcal{M}^{-1}$$

measured output responses

$$\delta Y = [\, \delta y^{(1)} \ \cdots \ \delta y^{(n)} \,]$$

A nonsingular linear map sends a basis to a basis.

Theorem (local probe informativeness)

Let $F_{\mathcal{G}}$ be C^2 near $\bar{y}$, and suppose

$$\mathcal{M} = DF_{\mathcal{G}}(\bar{y}) \quad \text{is nonsingular.}$$

Then $F_{\mathcal{G}}^{-1}$ exists locally, and the coordinate probes $\bar{w} + \kappa e_i$ yield

$$\delta Y = \kappa \mathcal{M}^{-1} + O(\kappa^2).$$

Consequently, there exists $\kappa_0 > 0$ such that

$$0 < |\kappa| < \kappa_0 \quad \Longrightarrow \quad \begin{aligned} &\det(\delta Y) \neq 0, \\ &\|\delta Y^{-1}\| = O(|\kappa|^{-1}). \end{aligned}$$

[Sharf, Z, TCNS 2023]

Tradeoff: small probes justify the local approximation, but amplify noise under inversion.

The Exception: An Invisible Common Mode

If the local node stiffness vanishes, $\mathcal{M}$ is Laplacian-like and $\mathcal{M}\mathbb{1} = 0$.

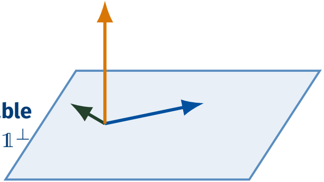

Work in the identifiable subspace

Use $n-1$ difference probes,

$$\delta w^{(i)} = \kappa(e_i - e_n), \qquad i = 1, \dots, n-1,$$

and project the measured responses onto $\mathbb{1}^\perp$. For small κ, these columns form a basis of that subspace.

Complete the coordinates with

$$\delta y^{(n)} = \kappa\mathbb{1}.$$

This last column is **synthetic**: it is appended algebraically, not produced by commanding an output.

Willems' Fundamental Lemma: A Shared Spanning View

Fundamental lemma: behavioral span

For a controllable LTI system with sufficiently rich measured data (u^d, y^d), every admissible length-L trajectory satisfies

$$\begin{bmatrix} \mathcal{H}_L(u^d) \\ \mathcal{H}_L(y^d) \end{bmatrix} \alpha = \begin{bmatrix} u^\star_{[0,L-1]} \\ y^\star_{[0,L-1]} \end{bmatrix}.$$

Here $\alpha \in \mathbb{R}^{T-L+1}$ weights the measured Hankel columns; those columns span all admissible finite trajectories.

This work: sensitivity span

For a regular equilibrium of a convergent nonlinear network, nearby constant probes produce a full-rank δY:

$$\widehat{\mathcal{M}}\,\delta Y \approx \delta W, \qquad \widehat{\mathcal{M}} = \delta W\,\delta Y^{-1}.$$

Its data columns span local equilibrium directions.

Analogy, not implication

This work does not apply the fundamental lemma. It borrows the spanning viewpoint, but changes the represented object: trajectories $\rightarrow$ local equilibrium Jacobian.

J.C. Willems, P. Rapisarda, I. Markovsky, and B.L.M. De Moor, "A note on persistency of excitation," *Systems & Control Letters*, 2005.

Network Identification Algorithm

A sequence of equilibrium queries becomes one structured matrix estimate.

1. Baseline query

$\bar{w}^{(0)} \mapsto \bar{y}^{(0)}$

wait for equilibrium

2. Local probes

$\bar{w}^{(0)} + \delta w^{(i)} \mapsto \bar{y}^{(i)}$

$i = 1, \ldots, n$

3. Assemble data

$\delta W = [\, \delta w^{(i)} \,]$

$\delta Y = [\, \bar{y}^{(i)} - \bar{y}^{(0)} \,]$

4. Estimate sensitivity

$\widehat{\mathcal{M}} = \delta W \, \delta Y^{-1}$

Experiment design

Use small, independent constant perturbations and sample only after settling.
The measured displacements must span the informative equilibrium directions.

Decode the off-diagonal entries: *for* $i \neq j$,

$$|\widehat{\mathcal{M}}_{ij}| > \varepsilon \Rightarrow \{i, j\} \in \widehat{\mathbb{E}}, \qquad \widehat{\nu}_{ij} = -\widehat{\mathcal{M}}_{ij}/d_{ij}, \quad d_{ij} = g'_{ij}(\bar{y}_i^{(0)} - \bar{y}_j^{(0)}).$$

Probe Size: Two Competing Errors

The same inverse that reconstructs the Jacobian also amplifies small errors.

Nonlinear finite-difference error

The Taylor remainder satisfies

$$\widehat{\mathcal{M}} - \mathcal{M} = R\delta Y^{-1}, \qquad \|R\| = O(\kappa^2).$$

Since $\|\delta Y^{-1}\| = O(\kappa^{-1})$,

$$\|R\delta Y^{-1}\| = O(\kappa).$$

Curvature bias scales as $O(\kappa)$.

Measurement and settling error

Finite waiting and sensor noise give

$$y_{\mathrm{meas}}(T) = \bar{y} + e_{\mathrm{settle}} + e_{\mathrm{sensor}},$$
$$\widetilde{\delta Y} = \delta Y + \Delta Y, \qquad \|\Delta Y\| \leq \eta.$$

For a small relative perturbation,

$$\|\Delta\widehat{\mathcal{M}}_{\mathrm{meas}}\| \lesssim \eta\|\delta Y^{-1}\| = O(\eta/\kappa)\,.$$

Measurement error scales as $O(\eta/\kappa)$.

Choosing Probe Size and the Error Threshold

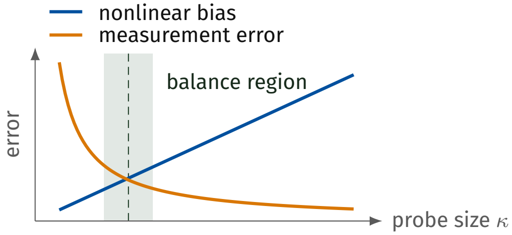

Schematic scaling from the error bounds.

Balancing the two terms suggests

$$\kappa \propto \sqrt{\eta},$$

up to problem-dependent constants.

Error scaling and threshold

Error scaling. For small probes,

$$|\widehat{\mathcal{M}}_{ij} - \mathcal{M}_{ij}| \lesssim C_{\mathrm{nl}}\kappa + C_{\mathrm{meas}}\frac{\eta}{\kappa}.$$

Large κ: nonlinear bias.

Small κ: noise amplification.

Thresholding. If $|\widehat{\mathcal{M}}_{ij} - \mathcal{M}_{ij}| \leq m$, use

$$\boxed{\varepsilon = 2m}.$$

Then every nonedge satisfies

$$|\widehat{\mathcal{M}}_{ij}| < \varepsilon.$$

Exact recovery also requires true edge entries to lie sufficiently far from zero.

When Does Thresholding Recover the Exact Graph?

Assume $|\widehat{\mathcal{M}}_{ij} - \mathcal{M}_{ij}| \leq m \quad (i \neq j), \qquad a_{\min} := \min_{\{i,j\}\in\mathbb{E}} |\nu_{ij} d_{ij}|.$

nonedge

$$\mathcal{M}_{ij} = 0$$
$$\Rightarrow |\widehat{\mathcal{M}}_{ij}| \leq m$$

threshold

$$\varepsilon = 2m(\kappa, \eta)$$

true edge

$$|\mathcal{M}_{ij}| \geq a_{\min}$$
$$\Rightarrow |\widehat{\mathcal{M}}_{ij}| \geq a_{\min} - m$$

Exact topology recovery

If $a_{\min} \geq 4m$, then

$|\widehat{\mathcal{M}}_{ij}| \leq m$ for nonedges, $\qquad |\widehat{\mathcal{M}}_{ij}| \geq 3m$ for true edges.

Hence the threshold $\varepsilon = 2m$ separates the two cases, so

$$\widehat{\mathbb{E}} = \mathbb{E}.$$

Why Worry About Complexity?

Identification has both a **physical experiment cost** and a **numerical reconstruction cost**.

Physical time

Each query must be held until the network settles.

wall clock $\approx rT_{\text{settle}}$

Too many probes can dominate the total duration.

Data size

The unknown structured Jacobian is $n \times n$, and the experiment matrix contains n^2 entries.

$\Theta(n^2)$ data entries

Even reading dense data has a quadratic cost.

Linear algebra

Reconstructing $\mathcal{M}$ requires solving an $n \times n$ matrix problem.

standard dense: $O(n^3)$

At large n, post-processing is no longer negligible.

Scalability questions: *How few settled experiments suffice? How much arithmetic is unavoidable? Does the algorithm attain both limits?*

Regular Case: Cost of the Proposed Method

Experimental cost

One baseline

$$\bar{w}^{(0)} \mapsto \bar{y}^{(0)}$$

plus one probe in each coordinate

$$\bar{w}^{(0)} + \kappa e_i \mapsto \bar{y}^{(i)}$$

$$i = 1, \ldots, n.$$

$$n + 1$$

settled measurements

Arithmetic cost: enforce structure first

1. Raw finite-difference estimate:

$$\widehat{\mathcal{M}}_{\mathrm{raw}} := \kappa \delta Y^{-1} \qquad (\delta W = \kappa I).$$

2. Its inverse is already in the data:

$$B := \widehat{\mathcal{M}}_{\mathrm{raw}}^{-1} = \kappa^{-1} \delta Y.$$

3. Enforce the known network structure.

Exactly, $B = \mathcal{M}^{-1}$ is symmetric positive definite (SPD). Small errors may add a skew part, so use

$$B_{\mathrm{sym}} := \tfrac{1}{2}(B + B^\top), \qquad \boxed{\widehat{\mathcal{M}}_{\mathrm{spd}} := B_{\mathrm{sym}}^{-1}}.$$

Boxed: a structure-enforced estimate, not an identity.

Cost: $O(n^{\omega_1})$, $\omega_1 :=$ exponent for SPD matrix inversion.

$2 \leq \omega_1 \leq \omega$, where ω is the unknown optimal exponent.

Theory: $\omega < 2.3728639$ [Le Gall, ISSAC'14]; dense: $O(n^3)$.

Complexity Lower Bounds and Optimality

Theorem (worst-case complexity lower bounds)

For the admissible class of n-node networks, any method guaranteed to recover the graph requires at least $n-1$ steady-state output measurements in the worst case.

Any deterministic method with this guarantee requires $\Omega(n^{\omega_1})$ arithmetic operations.

Regular-case comparison

	Unavoidable	**Proposed method**
Measurements	at least $n-1$	$n+1$
Operations	at least $\Omega(n^{\omega_1})$	$O(n^{\omega_1})$

Conclusion: the algorithm matches both unavoidable growth rates.

Example: 100 Oscillators With Dry Friction

Coupled mechanical model

$$\ddot{x}_i + F_i \operatorname{sgn}(\dot{x}_i) + a_i x_i = u_i + w_i, \qquad y_i = x_i,$$
$$u_i = -\sum_{j \in \mathcal{N}_i} \nu_{ij}(x_i - x_j).$$

The local springs a_i are known; the elastic coupling graph and weights ν_{ij} are unknown.

Equilibrium measurement

$$w_i = a_i \bar{y}_i + \sum_{j \in \mathcal{N}_i} \nu_{ij}(\bar{y}_i - \bar{y}_j).$$

Dry friction shapes the nonsmooth **transient**; the settled input–output equation is linear.

Network and probing experiment

- $n = 100$, Erdős–Rényi $p = 0.25$;
- $F_i \in [1, 10]$, $a_i \in [10^{-2}, 1]$;
- $\nu_{ij} \in [0.1, 10]$, log-uniform;
- probe scale $\kappa = 0.1$;
- threshold $\varepsilon = 0.01$, dwell time 100 s.

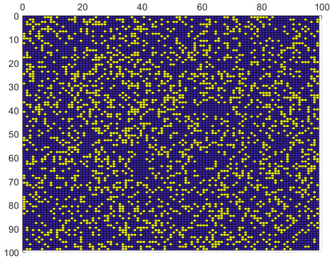

Interaction pattern recovered exactly

What Does a κ-Sweep Show Here?

For this linear equilibrium model, $\delta Y = \kappa \mathcal{M}^{-1}$ exactly. Exact settled data give the same $\mathcal{M}$ for every κ.

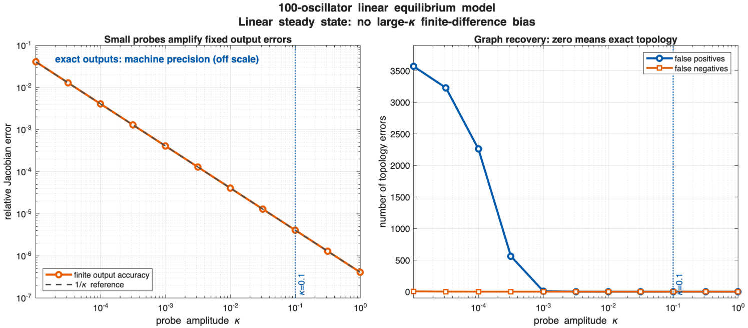

Left: *lower is better; fixed measurement or settling errors are amplified by $1/\kappa$.*
Right: *zero means exact graph recovery.*
Because this equilibrium equation is linear, large probes create no curvature penalty.

A Nonlinear Oscillator That Exhibits the Tradeoff

Hardening springs: $\beta_i > 0$

$$\ddot{x}_i + F_i \operatorname{sgn}(\dot{x}_i) + a_i x_i + \beta_i x_i^3 = u_i + w_i, \qquad y_i = x_i,$$
$$u_i = -\sum_{j \in \mathcal{N}_i} \nu_{ij}(x_i - x_j).$$

Its equilibrium equation follows directly

$$w_i = a_i y_i + \beta_i y_i^3 + \sum_{j \in \mathcal{N}_i} \nu_{ij}(y_i - y_j).$$

For $i \neq j$, $\partial w_i / \partial y_j = -\nu_{ij}$ on an edge and is zero otherwise: the off-diagonal Jacobian reveals the graph and weights.

Why probe size matters

Too small: the output change is comparable to sensing and settling error.

Too large: cubic curvature makes a finite secant differ from the local slope.

Simulation parameters

Same graph; $\beta_i \in [\sqrt{10}, 10]$, measurement noise $\sigma_y = 10^{-9}$, and a nonzero baseline.

Observed Probe-Amplitude Tradeoff

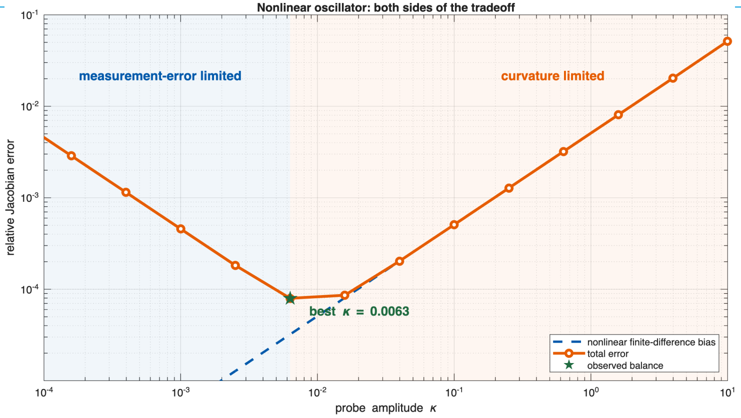

Observed balance and takeaway

$\kappa_\star \approx 6.3 \times 10^{-3}$ gives relative Jacobian error 7.96×10^{-5} and **exact graph recovery**. Balance output-error amplification against curvature.

Choosing κ Without Knowing the True Graph

For each candidate κ: fit $\widehat{\mathcal{M}}$ using the coordinate probes, then test it on new probe directions.

coordinate probes $\rightarrow$ $\widehat{\mathcal{M}}(\kappa)$ $\rightarrow$ predict new test probes $\rightarrow$ prediction residual

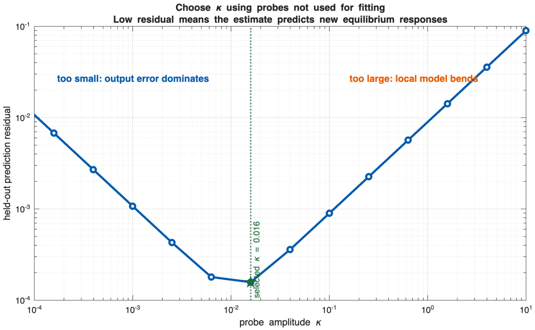

Observable tuning rule

Choose κ near the minimum held-out prediction residual.

In this example: $\kappa_{\text{val}} = 1.58 \times 10^{-2}, \quad \kappa_{\text{true}} = 6.31 \times 10^{-3}$, and both recover the exact graph.

Do We Really Need an Actuator at Every Node?

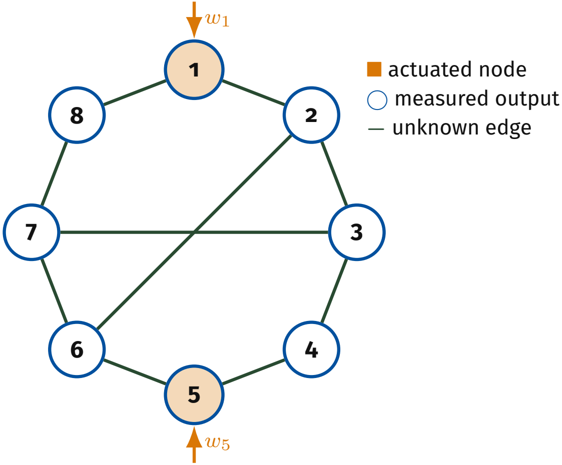

The obvious criticism

Coordinate probing appears to require an independent input at every node, together with measurement of every output.

A route beyond full actuation

That access pattern is **sufficient**, but need not be fundamental. Repeated equilibria and network structure can make restricted inputs informative. Here $\mathbb{V}_{\mathrm{probe}} = \{1,5\}$ while $\mathbb{V}_{\mathrm{meas}} = \mathbb{V}$: only two nodes are excited; all eight outputs are still measured.

Limited Probing: How Many Equilibria Are Needed?

Restricted experiment

$$k_i^{-1}(y_i) = a_i y_i + b_i y_i^3, \qquad g_{ij}(\zeta) = \nu_{ij}\zeta.$$

Only w_1, w_5 vary. No topology is assumed: 28 candidates, 10 true edges. Data: $S = 11$ equilibria (one baseline plus ten perturbations).

Structural minimum

At equilibrium s, the data reveal $Ly^{(s)}$. The unknown K_8 Laplacian also satisfies $L\mathbb{1} = 0$.

Hence L is determined only if

$$\mathrm{rank}[\,\mathbb{1}\ y^{(0)}\ \cdots\ y^{(S-1)}\,] = 8.$$

Therefore $S \geq 7$: one baseline plus at least six additional equilibria.

Why use eleven here?

$S = 11$ is **not a minimum**. Four extra operating points overdetermine the regression and make it easier to obtain a full-rank, better-conditioned dataset from only two inputs.

How Can Two Inputs Reveal 28 Candidate Edges?

Use every equilibrium

At equilibrium s, use the network balance

$$\underbrace{w^{(s)} - k^{-1}(y^{(s)})}_{\text{network contribution}} = \underbrace{\Phi(y^{(s)})\nu}_{\text{unknown edge weights}} .$$

Stack the eleven points:

$$\mathbf{r} = \boldsymbol{\Phi}\nu, \qquad \boldsymbol{\Phi} = \begin{bmatrix} \Phi(y^{(0)}) \\ \Phi(y^{(1)}) \\ \vdots \\ \Phi(y^{(10)}) \end{bmatrix} \in \mathbb{R}^{88\times 28}.$$

Nonlinearity varies $\Phi(y^{(s)})$ across operating points.

Why two input channels can suffice

The nonlinear probe-to-output map bends as the operating point changes. Its sampled outputs can therefore span more directions than the two-dimensional input space.

The decisive numerical test

$$\boxed{\operatorname{rank} \boldsymbol{\Phi} = 28}.$$

All 28 candidate weights are uniquely determined in the noiseless experiment.

Proof of Possibility: Two Actuators Recover the Graph

Recovered graph: edge labels are $\widehat{\nu}_{ij}$

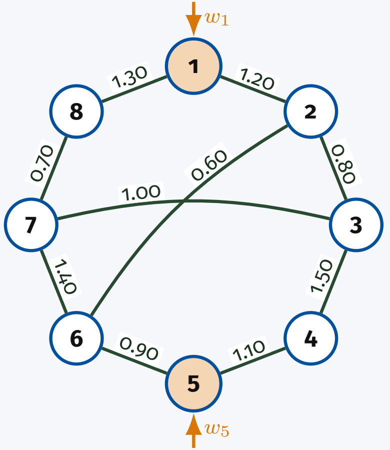

What the data certify

- access: actuate nodes $1, 5$; measure all 8 outputs;
- data: $S = 11$ equilibria;
- regression: 28 candidates and $\operatorname{rank} \mathbf{\Phi} = 28$;
- support: $10/10$ edges recovered, $0/18$ false positives;
- max. weight error: 2.4×10^{-9}.

What the Example Establishes—and What Comes Next

It establishes a possibility

- full actuation is sufficient, not necessary;
- nonlinear operating-point variation can compensate for missing input directions;
- structure-aware regression can recover the complete graph from restricted actuation.

It does not yet remove every limitation

- all node outputs are still measured;
- the demonstration is noiseless;
- full rank alone does not guarantee good conditioning.

Next: *choose actuator locations and operating points to maximize $\sigma_{\min}(\mathbf{\Phi})$, then ask which output measurements can also be removed.*

Concluding Remarks

Learn the structured equilibrium Jacobian, not the nonlinear transient dynamics.

Representation

Global convergence turns the network into a repeatable steady-state input-output oracle.

Structure

The local Jacobian is diagonal node physics plus a weighted graph Laplacian.

Data

Rich equilibrium probes reveal topology and weights with nearly minimal effort.

Data-driven connection: *experiment design builds a spanning representation, but of local equilibrium sensitivity rather than trajectory behavior.*

References

- M. Sharf and D. Zelazo, "Network Identification for Diffusively Coupled Networks With Minimal Time Complexity," *IEEE Transactions on Control of Network Systems*, 10(3):1616–1628, 2023.
- M. Sharf and D. Zelazo, "Monitoring Link Faults in Nonlinear Diffusively-Coupled Networks," *IEEE Transactions on Automatic Control*, 67(6):2857–2872, 2022. Source for convergence profiles and comparison-ODE sampling times.
- M. Sharf and D. Zelazo, "Network Identification: A Passivity and Network Optimization Approach," *IEEE Conference on Decision and Control*, 2018.
- M. Bürger, D. Zelazo, and F. Allgöwer, "Duality and Network Theory in Passivity-Based Cooperative Control," *Automatica*, 50(8):2051–2061, 2014.