What Rigidity Hides: Controllability and Transmission in Distance-Based Formations

Rigidity Workshop - Lancaster University

Solomon Goldgraber Casspi and **Daniel Zelazo**

The Stephen B. Klein Faculty of Aerospace Engineering, Technion

September 2026

Formation control

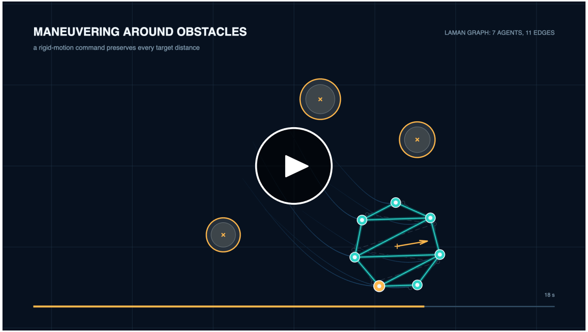

Formation control coordinates agents through local interactions:

- acquire a prescribed geometric shape
- preserve the shape while moving

Applications

- Aerial robots: mapping and inspection
- Satellite formations: coordinated sensing
- Mobile robots: transport and coverage

A simple feedback model

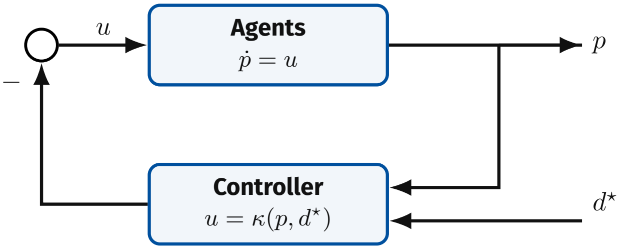

Agent model

For n agents,

$$p = (p_1, \ldots, p_n), \qquad p_k \in \mathbb{R}^d,$$

and we use the simple dynamics

$$\dot{p}_k = u_k.$$

Control language

p : ***state*** $\qquad u$: ***control input***

The input is continually recomputed from measurements of the evolving state: ***feedback****.*

Formation control as gradient descent

Rigidity map

$$\boxed{F_{\mathcal{G}}(p) = \left(\|p_i - p_j\|^2\right)_{(i,j)\in\mathcal{E}}}$$

A target shape $p^\star$ specifies the desired value $F_{\mathcal{G}}(p^\star)$.

Distance-error potential

$$V(p) = \frac{1}{4}\left\|F_{\mathcal{G}}(p) - F_{\mathcal{G}}(p^\star)\right\|^2.$$

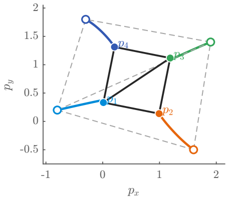

Gradient feedback

$$\boxed{\begin{aligned}\dot{p} &= u = -\nabla V(p),\\ &= -R(p)^\top\left(F_{\mathcal{G}}(p) - F_{\mathcal{G}}(p^\star)\right).\end{aligned}}$$

a **distributed** implementation

$$\dot{p}_i = \sum_{ij\in\mathcal{E}}\left(\|p_i - p_j\|^2 - d_{ij}^\star\right)(p_j - p_i)$$

Gradient-flow structure

- $\dot{V} = -\|\nabla V(p)\|^2 \leq 0$
- trajectories approach critical set of V
- infinitesimally rigid targets are locally attracting, modulo trivial motions

Krick, Broucke & Francis, *IJRNC*, 2009.

Beyond stability: a control-theoretic viewpoint

Control theorists are also interested in how localized inputs and disturbances **propagate through the dynamics**, and what can be inferred from **partial measurements**.

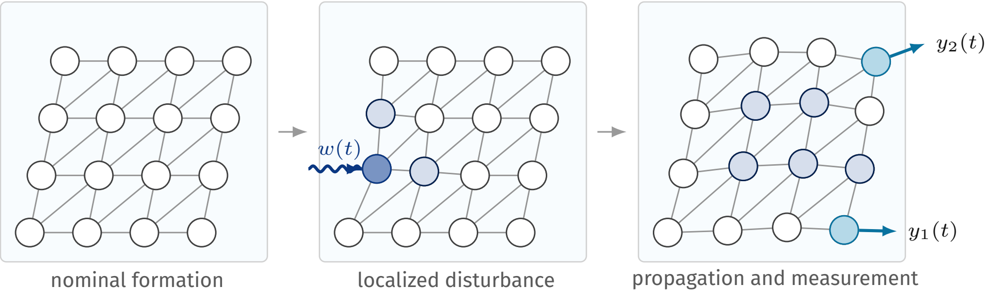

Which motions can be excited?	Which motions can be observed?	How is an input transmitted?
Controllability	Observability	Transmission

The story so far

1. Controllability of Distance-Based Formations

2. Transmission in Distance-Based Formations

3. Nonlinear Steering

Localized inputs to the formation

To study how external signals affect the formation, we add an input acting on a single agent i:

Input-driven formation dynamics

$$\dot{p} = -R(p)^\top \big(F_{\mathcal{G}}(p) - F_{\mathcal{G}}(p^\star)\big) + B_i w$$

Here

$$w \in \mathbb{R}^d, \quad B_i = \mathrm{e}_i \otimes I_d,$$

so the external input acts only on agent i.

B_i specifies **where the signal enters** the formation.

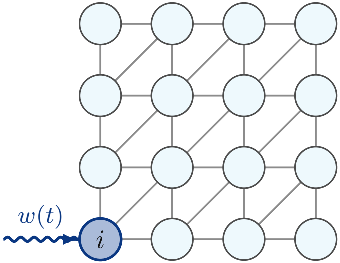

A disturbance or control signal is applied at a **localized subset** of the formation.

Local dynamics around the target formation

We are interested in the response to **small disturbances** when the formation is operating near its desired shape.

Nonlinear formation dynamics

$$\dot{p} = -\nabla V(p) + B_i w$$

linearize about $p^\star$

$$x = p - p^\star$$

Local input-driven dynamics

$$\dot{x} = -\nabla^2 V(p^\star)x + B_i w$$

At the desired formation,

$$\nabla^2 V(p^\star) = R(p^\star)^\top R(p^\star).$$

$$\Longrightarrow \quad \boxed{\dot{x} = -R(p^\star)^\top R(p^\star)x + B_i w}$$

The key point: the local system matrix is determined entirely by the **geometry of the framework**.

Controllability: which motions can an input generate?

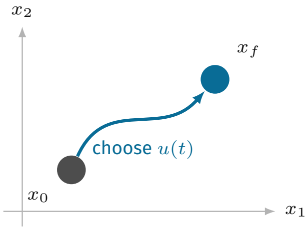

A system is **controllable** if the input can steer the state from any x_0 to any x_f in finite time.

The input enters through B, and the dynamics spread its influence through

$$B,\ AB,\ A^2B,\ \ldots$$

(matrix exponentials, Cayley-Hamilton)

For the linear system

$$\dot{x} = Ax + Bu,$$

controllability can be checked by each mode.

PBH controllability test
The pair (A, B) is controllable if and only if

$$B^\top v \neq 0$$

for every eigenvector v of A, i.e., for every nonzero v satisfying $Av = \lambda v$.

Controllability: which motions can an input generate?

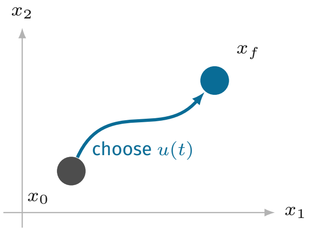

A system is **controllable** if the input can steer the state from any x_0 to any x_f in finite time.

The input enters through B, and the dynamics spread its influence through

$$B,\ AB,\ A^2B,\ \ldots$$

For the linear system

$$\dot{x} = Ax + Bu,$$

controllability can be checked by each mode.

PBH controllability test

The pair (A, B) is controllable if and only if

$$B^\top v \neq 0$$

for every eigenvector v of A, i.e., for every nonzero v satisfying $Av = \lambda v$.

Uncontrollable subspace $\bar{C} := \mathrm{span}\{$*uncontrollable eigenmodes*$\}$

Controllability: which motions can an input generate?

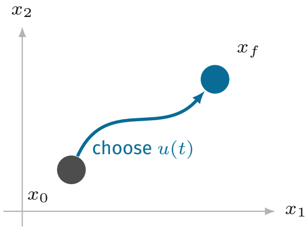

A system is **controllable** if the input can steer the state from any x_0 to any x_f in finite time.

The input enters through B, and the dynamics spread its influence through

$$B,\ AB,\ A^2B,\ \ldots$$

For the linear system

$$\dot{x} = Ax + Bu,$$

controllability can be checked by each mode.

PBH controllability test

The pair (A, B) is controllable if and only if

$$B^\top v \neq 0$$

for every eigenvector v of A, i.e., for every nonzero v satisfying $Av = \lambda v$.

$$\bar{\mathcal{C}} = \underbrace{\ker R(p^\star) \cap \ker B^\top}_{\textit{\textbf{Hidden RBMs}}\ (\lambda=0)} \oplus \bigoplus_{\lambda<0} \underbrace{\left(\ker(-R(p^\star)^\top R(p^\star) - \lambda I) \cap \ker B^\top\right)}_{\textit{\textbf{Hidden deformations}}}.$$

Dimension count

For a localized input at agent i, $B_i^\top v = v_i$. Thus PBH becomes a geometric question:

Is there an eigenmode that leaves agent i stationary?

The constraint $v_i = 0$ imposes only d scalar constraints on the rigid-body motion space.

Therefore

$$\dim\big(\ker R(p^\star) \cap \ker B_i^\top\big) \geq \frac{d(d+1)}{2} - d = \boxed{\frac{d(d-1)}{2}}.$$

- $d = 2 \Rightarrow$ at least one hidden rigid-body mode.
- $d = 3 \Rightarrow$ at least three hidden rigid-body modes.

What do these hidden rigid-body modes look like?

What does the pinning constraint mean?

PBH gave the condition for a hidden rigid-body mode: $\boxed{v_i = 0}$.
Every infinitesimal trivial motion has the form

$$v_j = t + \Omega p_j, \qquad \Omega^\top = -\Omega$$

t is a translation; Ω is an infinitesimal rotation.
At the actuated agent,

$$0 = v_i = t + \Omega p_i \Rightarrow t = -\Omega p_i.$$

Substituting back,

$$\boxed{v_j = \Omega(p_j - p_i).}$$

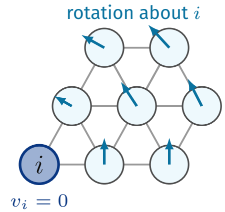

The only rigid-body motions fixing agent i are rotations about i.

This explains the unavoidable hidden rigid-body modes.
But are there also hidden **shape-changing** modes?

What remains unexplained?

From the modal splitting,

$$\bar{\mathcal{C}} = \underbrace{\ker R(p^\star) \cap \ker B_i^\top}_{\text{hidden rigid-body modes}} \oplus \underbrace{\bigoplus_{\lambda<0} \left(E_\lambda \cap \ker B_i^\top\right)}_{\text{hidden deformation modes}} .$$

We have now characterized the first term:

$$\ker R(p^\star) \cap \ker B_i^\top = \{\text{infinitesimal rotations about } i\}.$$

The remaining question:
What geometric structure constrains the **shape-changing uncontrollable modes**?

The local rotational subspace

To constrain hidden deformation modes, consider motions that are locally invisible to the star of agent i:

$$v_i = 0, \qquad (p_j^\star - p_i^\star)^\top v_j = 0 \quad \forall (i,j) \in \mathcal{E}.$$

Local rotational subspace

$$\mathcal{T}_i = \left\{ v : \; v_i = 0, \; (p_j^\star - p_i^\star)^\top v_j = 0 \; \forall (i,j) \in \mathcal{E} \right\}.$$

So each neighbor may move **tangentially** around the fixed agent i, without changing the length of edge (i,j) to first order.

Dimension: pin i ($-d$), constrain each neighbor (-1):

$$\dim \mathcal{T}_i = d(n-1) - \deg(i).$$

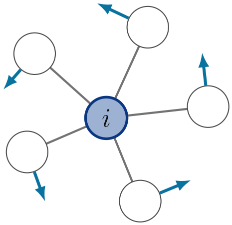

tangent to incident edges

The local rotational subspace

To constrain hidden deformation modes, consider motions that are locally invisible to the star of agent i:

$$v_i = 0, \qquad (p_j^\star - p_i^\star)^\top v_j = 0 \quad \forall (i,j) \in \mathcal{E}.$$

Local rotational subspace

$$\mathcal{T}_i = \left\{ v : \; v_i = 0, \; (p_j^\star - p_i^\star)^\top v_j = 0 \; \forall (i,j) \in \mathcal{E} \right\}.$$

So each neighbor may move **tangentially** around the fixed agent i, without changing the length of edge (i,j) to first order.

Dimension: pin i ($-d$), constrain each neighbor (-1):

$$\dim \mathcal{T}_i = d(n-1) - \deg(i).$$

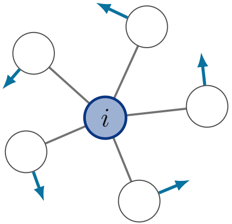

tangent to incident edges

$\mathcal{T}_i$ describes motions that the local star at i cannot see to first order.

Why uncontrollable modes lie in $\mathcal{T}_i$

For an uncontrollable eigenmode,

$$Av = \lambda v, \qquad B_i^\top v = 0 \quad \Rightarrow \quad v_i = 0, \qquad (Av)_i = 0.$$

With $A = R(p^\star)^\top R(p^\star)$,

$$(Av)_i = \sum_{j \in \mathcal{N}_i} z_{ij}\alpha_j = 0, \qquad z_{ij} := p_j^\star - p_i^\star, \qquad \alpha_j := z_{ij}^\top v_j.$$

Minimal-degree actuator. If $\deg(i) = d$, then the incident directions $\{z_{ij}\}_{j \in \mathcal{N}_i}$ are linearly independent. Hence

$$\sum_j z_{ij}\alpha_j = 0 \quad \Rightarrow \quad \alpha_j = 0 \quad \forall j.$$

Why uncontrollable modes lie in $\mathcal{T}_i$

Since $\alpha_j = z_{ij}^\top v_j$, we obtain

$$z_{ij}^\top v_j = 0 \qquad \forall j \in \mathcal{N}_i.$$

Together with $v_i = 0$, this is exactly the definition of $\mathcal{T}_i$. Therefore

$$\boxed{\bar{\mathcal{C}} \subseteq \mathcal{T}_i.}$$

Interpretation. Every uncontrollable mode is locally invisible at the actuated agent: agent i is fixed, and all incident edge lengths are unchanged to first order.

Why uncontrollable modes lie in $\mathcal{T}_i$

$\boxed{\bar{\mathcal{C}} \subseteq \mathcal{T}_i}$ but generally $\boxed{\bar{\mathcal{C}} \neq \mathcal{T}_i}$.

Example: a triangle actuated at agent 1

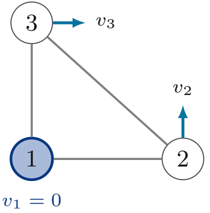

$$v_1 = (0,0), \qquad v_2 = (0,1), \qquad v_3 = (1,0).$$

Since $v_2 \perp (p_2 - p_1), \quad v_3 \perp (p_3 - p_1)$, we have

$$\boxed{v \in \mathcal{T}_1.}$$

But the motion changes the shape.
The edge opposite the actuator changes length to first order:

$$(p_2 - p_3)^\top (v_2 - v_3) = -2 \neq 0.$$

Moreover,

$$Av = -R^\top R v \notin \operatorname{span}\{v\},$$

so v is **not an eigenmode** of the formation dynamics.

Thus v is locally invisible at agent 1, but it is not an uncontrollable mode.

Shape recovery after a localized impulse

A localized impulse gives agent i a kick:

$$w(t) = w_0 \delta(t).$$

After deformation modes decay, only rigid-body motion remains:

$$\delta p(\infty) = c_x v_x + c_y v_y + c_r v_r.$$

The rotational coefficient is local:

$$\boxed{c_r = \langle [v_r]_i, w_0 \rangle.}$$

Recovery criterion

$$\langle [v_r]_i, w_0 \rangle = 0 \quad \iff \quad c_r = 0.$$

The asymptotic motion is pure translation.

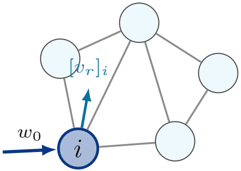

recovery depends on the angle between w_0 and $[v_r]_i$

Geometric message

If the kick is orthogonal to the local rotational velocity $[v_r]_i$, it does not excite rotation.

If not, a rotational component remains.

Shape recovery depends on the impulse direction

For a planar formation,

$$c_r = \langle [v_r]_i, w_0 \rangle$$

determines whether the impulse excites the rotational mode.

Orthogonal impulse

$$\langle [v_r]_i, w_0 \rangle = 0 \quad \Rightarrow \quad c_r = 0.$$

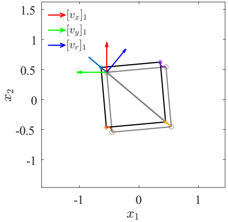

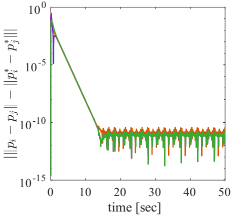

The edge errors decay to zero: **the formation recovers its shape.**

Shape recovery depends on the impulse direction

For a planar formation,

$$c_r = \langle [v_r]_i, w_0 \rangle$$

determines whether the impulse excites the rotational mode.

Impulse excites rotation

$$\langle [v_r]_i, w_0 \rangle \neq 0 \quad \Rightarrow \quad c_r \neq 0.$$

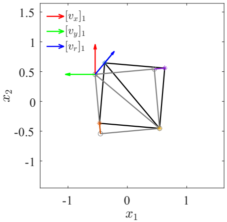

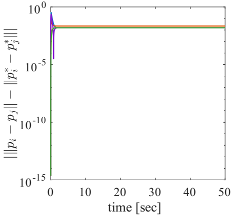

The edge errors converge to nonzero values: **the formation does not recover its original shape.**

Part I: what we found

Hidden modes depend on the full graph and configuration, but are **bounded** by local geometry.

(i) Hidden modes via PBH

A localized actuator misses every eigenmode that leaves agent i stationary: hidden rigid-body motions and hidden deformation modes.

(ii) Local outer bound $\mathcal{T}_i$

Every hidden mode lies in $\mathcal{T}_i$, which depends only on the star of i. The full graph determines *which* modes within $\mathcal{T}_i$ are actually hidden.

(iii) Shape-recovery criterion

After an impulse w_0 at agent i, the formation recovers its shape if and only if $w_0 \perp [v_r]_i$.

$\dim \mathcal{T}_i = d(n-1) - \deg(i)$: *higher degree at i shrinks the bound — no global information needed.*

From controllability to transmission

So far we asked:

Which internal modes can a localized input excite?

But often the relevant question is not the full internal state.

If a disturbance enters at agent i,
what can be detected at another agent j?

This leads to an input-output model:

$$\dot{x} = Ax + B_i w, \qquad y = C_j x.$$

Transmission asks how signals propagate from input to output.

The story so far

1. Controllability of Distance-Based Formations

2. Transmission in Distance-Based Formations

3. Nonlinear Steering

The input-output channel

Question If node i is disturbed by $w(t)$, what is the position deviation δp_j of node j?

The Laplace transform of the linearised dynamics gives a single matrix function encoding the complete channel:

$$Y(s) = \underbrace{C_j(sI - A)^{-1}B_i}_{G_{ji}(s)\,\in\,\mathbb{R}^{d\times d}}\, W(s).$$

Why transfer matrices?

$G_{ji}(s)$ captures the channel *frequency by frequency*:

- **poles** — natural modes of the formation,
- **zeros** — frequencies blocked entirely from node j, whatever the magnitude of w.

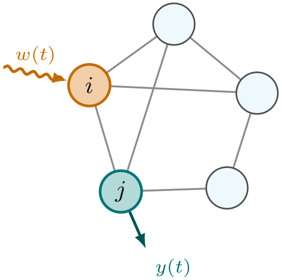

Node i (orange) inputs disturbance w; node j (teal) shows position deviation $y = \delta p_j$.

Transmission zeros

Transfer function from node i to node j:

$$G_{ji}(s) = C_j(sI + R^\top R)^{-1}B_i.$$

Modal decomposition via eigenmodes of $R^\top R$:

$$G_{ji}(s) = \underbrace{\frac{[P_0]_{ji}}{s}}_{\text{DC / RBM}} + \sum_{\lambda_k>0} \frac{[\varphi_k]_j[\varphi_k]_i^\top}{s+\lambda_k}.$$

Zero-frequency residue (ZFR) matrix

$$[P_0]_{ji} := \lim_{s\to 0} s\,G_{ji}(s).$$

A **steady-state zero** occurs iff $\det\big([P_0]_{ji}\big) = 0.$

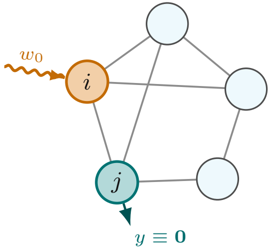

Sustained disturbance w_0 at i: if $\det([P_0]_{ji}) = 0$, node j stays stationary.

The ZFR matrix for rigid formations

For an **infinitesimally rigid** formation the ZFR matrix has the explicit form:

$$[P_0]_{ji} = \tfrac{1}{n}\, I_2 + [v_r]_j\, [v_r]_i^\top,$$

where $[v_r]_k = \frac{1}{\sqrt{J_p}}\,\Omega(p_k^* - p_{\mathrm{cm}})$ is the **rotation-mode component** at node k ($J_p = \sum_k \|p_k^* - p_{\mathrm{cm}}\|^2$, Ω a 90° rotation).

By **Sylvester's determinant identity** ($\det(I + uv^\top) = 1 + v^\top u$):

$$\det\big([P_0]_{ji}\big) = \tfrac{1}{n^2}\big(1 + n\,\langle [v_r]_i,\, [v_r]_j\rangle\big).$$

Rank-drop condition for rigid formations

G_{ji} has a **steady-state zero** $\quad\Longleftrightarrow\quad 1 + n\,\langle [v_r]_i,\, [v_r]_j\rangle = 0.$

Geometric reading: $[v_r]_k$ is tangent to a circle around p_{cm}. A zero arises when the inner product of the rotation-mode components at i and j equals $-1/n$.

Zero condition in formation systems

Substitute $[v_r]_k = \frac{1}{\sqrt{J_p}}\,\Omega(p_k^* - p_{\rm cm})$ into the rank-drop condition $1 + n\langle [v_r]_i, [v_r]_j\rangle = 0$:

$$1 + \tfrac{n}{J_p}\left\langle \Omega(p_i^* - p_{\rm cm}),\ \Omega(p_j^* - p_{\rm cm})\right\rangle = 0.$$

Since $\Omega^\top\Omega = I_2$:

$$\left\langle p_i^* - p_{\rm cm},\ p_j^* - p_{\rm cm}\right\rangle = -\tfrac{J_p}{n}.$$

Theorem 1 — geometric locus

G_{ji} has a **steady-state transmission zero** $\iff$ p_j^* lies on the affine line

$$\mathcal{L}_i = \left\{ x \in \mathbb{R}^2 \;\middle|\; \left\langle p_i^* - p_{\rm cm},\ x - p_{\rm cm}\right\rangle = -\tfrac{J_p}{n} \right\}.$$

$\mathcal{L}_i$ is **perpendicular** to $(p_i^* - p_{\rm cm})$ and lies on the *opposite side* of $p_{\rm cm}$ from node i.

Flexible formations: zeros are non-generic

For a **flexible** formation, the ZFR matrix $[P_0]_{ji}(p^*)$ depends on flexible-mode eigenvectors — complicated nonlinear functions of the configuration $p^* \in \mathbb{R}^{2n}$.

Proposition 1

For a generic connected flexible planar formation, the set of configurations $p^* \in \mathbb{R}^{2n}$ for which G_{ji} has a steady-state zero is a **real algebraic variety** of Lebesgue measure zero.

The polynomial $\det([P_0]_{ji}(p^*))$ is not identically zero: it vanishes only on a *proper closed* subset of $\mathbb{R}^{2n}$ — a lower-dimensional surface in configuration space.

A flexible formation at a ***generic*** *configuration has no steady-state zeros. Zeros are structurally exceptional: they cannot be read off the graph, only from specific geometry.*

Rigid formations: the variety collapses to a line

For a **rigid** formation, the Sylvester formula pins down the determinant explicitly:

$$\det\big([P_0]_{ji}\big) = \tfrac{1}{n^2}\big(1 + n\langle [v_r]_i,\, [v_r]_j\rangle\big).$$

Key observation. $[v_r]_j = \frac{1}{\sqrt{J_p}}\,\Omega(p_j^* - p_{\mathrm{cm}})$ is *linear* in p_j^*, so the zero condition is a **linear** equation in p_j^* — not a general polynomial. The algebraic variety drops to its lowest possible dimension.

The zero condition is *linear* in p_j^*, so the algebraic variety collapses to a single **affine line** $\mathcal{L}_i$ in $\mathbb{R}^2$ — Theorem 1, with an explicit formula shown on the next slide.

The geometric locus $\mathcal{L}_i$

The locus $\mathcal{L}_i$ is fully determined by the equilibrium configuration:

$$\mathcal{L}_i = \left\{ x \in \mathbb{R}^2 \;\middle|\; \langle p_i^* - p_{\mathrm{cm}},\, x - p_{\mathrm{cm}} \rangle = -\frac{J_p}{n} \right\}.$$

Geometric properties of $\mathcal{L}_i$:

- A **line** in $\mathbb{R}^2$
- $\perp$ to the vector $(p_i^* - p_{\mathrm{cm}})$
- On the *opposite* side of p_{cm} from node i
- Distance from p_{cm}: $\dfrac{J_p}{n\,\|p_i^* - p_{\mathrm{cm}}\|}$

Node $j \in \mathcal{L}_i$: a **blind node** — steady-state signal from i is **blocked**.

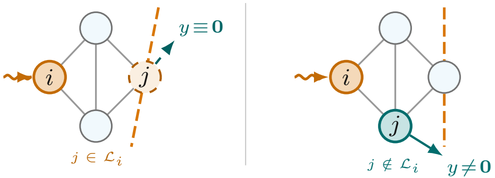

Dashed orange: locus $\mathcal{L}_i$.
Left: $j \in \mathcal{L}_i$ — signal blocked.
Right: $j \notin \mathcal{L}_i$ — signal reaches j.

The global transmission polygon $\mathcal{C}$

Theorem 1 constrains sensor placement for a *known* actuator. In practice, disturbances may enter at **any** node. The **admissible sensor set** must avoid every locus:

$$\mathcal{A} = \bigcap_{i \in \mathcal{V}} \left(\mathbb{R}^2 \setminus \mathcal{L}_i\right).$$

For each actuator, safe half-space containing p_{cm}:

$$\mathcal{H}_i = \left\{ x \in \mathbb{R}^2 \;\middle|\; \langle p_i^* - p_{\mathrm{cm}},\, x - p_{\mathrm{cm}} \rangle > -\frac{J_p}{n} \right\}.$$

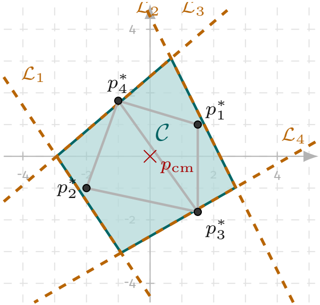

Global Transmission Polygon

$$\mathcal{C} := \bigcap_{i \in \mathcal{V}} \mathcal{H}_i \subseteq \mathcal{A}.$$

$\mathcal{C}$ is **convex**, non-empty, and gives a constructive sensor placement rule.

Proposition 2

If $p_j^* \in \mathcal{C}$, then G_{ji} has no steady-state zero for **any** $i \in \mathcal{V}$.

Part II: what we found

Transmission zeros are the formation-theoretic obstruction to signal propagation, shaped entirely by the framework geometry.

(i) Steady-state zeros via ZFR
$\det([P_0]_{ji}) = 0$: a disturbance direction at i leaves node j stationary ($y \equiv \mathbf{0}$).

(ii) Geometric locus $\mathcal{L}_i$
The set of positions creating a zero is an **affine line** in the configuration space.

(iii) Global polygon $\mathcal{C} = \bigcap_i \mathcal{H}_i$
Any sensor inside $\mathcal{C}$ receives every signal from **any** actuator with no blockage.

Each $\mathcal{L}_i$ is a line; any sensor there receives a zero from actuator i.

- ***Symmetric** formations: symmetry forces more agent pairs into the perpendicularity condition — loci land on other agents, shrinking $\mathcal{C}$.*
- ***Generic** formations: loci in general position, $\mathcal{C}$ stays large and placement is easy.*

The story so far

1. Controllability of Distance-Based Formations

2. Transmission in Distance-Based Formations

3. Nonlinear Steering

Nonlinear extension: from a tangent space to a manifold

The linearized analysis describes perturbations in $T_{p^*}\mathcal{M}$ around p^* — a **first-order approximation**. The true state evolves on the constraint manifold $\mathcal{M} = F^{-1}(F(p^*))$:

$$\dot{p} = \underbrace{f_0(p)}_{-R(p)^\top R(p)\, p} + \sum_{\alpha} u_\alpha \underbrace{B_i^\alpha}_{e_i \otimes e_\alpha}.$$

The drift f_0 encodes formation rigidity; the controls B_i^α are constant, but their brackets with f_0 can generate **directions not visible in the linearization**.

Key question: Can the **curvature** of $\mathcal{M}$ generate directions blocked by the transmission zeros of Part II?

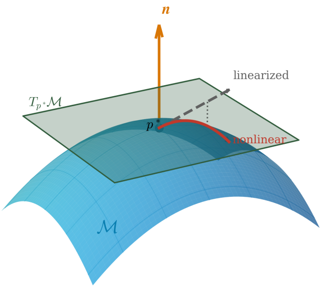

What the nonlinear analysis adds

Lie bracket: net displacement from a small input loop — brackets with the formation drift f_0 encode the curvature of $\mathcal{M}$.

Part II result (linear picture)

- Steady-state zeros: $\det([P_0]_{ji}) = 0$
- Channel $i \to j$ **blocked** when $p_j^* \in \mathcal{L}_i$
- DC disturbance at i leaves node j stationary

Nonlinear picture (on $\mathcal{M}$)

- Rigidity curvature encoded in $[f_0, B_i^\alpha]$
- Mixed brackets $[B_i^\alpha, [f_0, B_j^\beta]]$ couple curvature to independent axes
- May **open channels** blocked by ZFR zeros

Conclusion

Transmission zeros of Part II are an *obstruction in the linear picture*. Nonlinear brackets generated by the formation drift $f_0(p) = -R(p)^\top R(p)\, p$ may recover blocked directions — accessibility at generic equilibria *(in preparation)*.

Brackets recover hidden rotations

Degree-one brackets $\{B_i^\alpha, [f_0, B_i^\alpha], \ldots\}$ give exactly the **linear controllable subspace** $\mathcal{C}_i$.

Mixed brackets couple the rigidity curvature to *independent* input axes:

$$[B_i^\beta, [f_0, B_i^\alpha]](p^*), \quad \alpha \neq \beta.$$

Their projections span the missing **rotational space** $\mathcal{R}_i(p^*)$ at a generic equilibrium, so $\mathcal{D}(p^*) = T_{p^*}\mathcal{M}$.

Conjecture 1 — Accessibility (in preparation)

The single-actuator nonlinear formation system is **locally accessible** at almost all regular equilibria, even though its linearization is uncontrollable.

$$\text{Linear: } \mathcal{C}_i \subsetneq T_{p^*}\mathcal{M} \xrightarrow{\text{mixed brackets } [B_i^\beta, [f_0, B_i^\alpha]]} \text{Nonlinear: } \mathcal{D}(p^*) = T_{p^*}\mathcal{M}$$

No direction is permanently blocked — but accessibility does not say how fast or by what path.

Small-time controllability: an isotropy condition

Accessibility (Conj. 1): every direction is reachable via *some* sequence of inputs — no direction is permanently blocked.
STLC is stronger: every nearby state reachable in **small time**.
Bracket directions must be **reversible**.

The one-way drift obstruction

A small commutator loop produces a second-order rotation

$$\Delta\theta \;\propto\; u^\top [L(p^*)]_{ii}\Omega u.$$

But this term is quadratic in the input direction:

$$\Delta\theta(-u) = \Delta\theta(u).$$

Thus reversing the localized input does *not* reverse the induced rotation—the key obstruction to small-time local controllability.

Conjecture 2 — Isotropy rule *(in preparation)*

One-way drift vanishes $\iff$ **isotropically placed** neighbours: $[L(p^*)]_{ii} = \lambda I_d$. *Only node i's local neighbourhood decides this.*

Left: *neighbours isotropic (λI_d)*
$\Rightarrow$ *drift cancels, **STLC holds**.*
Right: *elongated neighbourhood*
$\Rightarrow$ *one-way drift, **STLC fails**.*

Part III: what we found

(i) Accessibility (Conjecture 1)
Mixed brackets $[B_i^\beta, [f_0, B_i^\alpha]]$, $\alpha \neq \beta$, span the missing rotational space at a generic equilibrium: the system is **locally accessible**.

(ii) One-way rotational drift
Repeating the same input in the commutator loop creates an irreversible rotation that **blocks STLC**.

(iii) Isotropy rule (Conjecture 2)
STLC obstruction vanishes iff $[L(p^*)]_{ii} = \lambda I_d$: **neighborhood geometry** at node i decides this.

***Accessibility**: no direction is permanently blocked by curvature.*

***STLC**: reaching those directions in small time requires **local isotropy**.*

Conclusion

Three perspectives on one rigidity-theoretic formation system:

Part I — Controllability. A single actuator at node i cannot drive the state into the rotational subspace $T_i(p^*)$; missing directions are permanent in the linear picture.

Part II — Transmission zeros. The ZFR matrix $[P_0]_{ji}$ reveals which neighbors are frozen; the locus $\mathcal{L}_i$ and polygon $\mathcal{C}$ characterise the obstruction geometrically.

Part III — Nonlinear steering. Mixed brackets $[B_i^\beta, [f_0, B_i^\alpha]]$ recover the missing space; the STLC obstruction vanishes iff $[L(p^*)]_{ii} = \lambda I_d$.

Open questions

Theory: prove Conjectures 1 & 2.

Algorithms: practical bracket-based steering manoeuvres.

Geometry: extend to $\mathbb{R}^3$ (SO(3) modes).

Directed sensing: how does the obstruction geometry change?

New paper: directed formation control via edge dynamics

Formation control from the generic combinatorial viewpoint: edge dynamics and directed sensing

Theran Zelazo Dewar Schulze • arXiv:2609.13423

Agent i senses only incoming edges; the directed controller is: $$u_i = -\sum_{j:(j,i)\in\vec{\mathcal{E}}} w_{ij}\left(\|p_i - p_j\|^2 - d_{ij}^{*2}\right)(p_i - p_j)$$

Admissibility

Admissibility is a generic spectral property of a compatible controller. Intuitively, it says the controller has no zero modes in the physically realizable edge directions, making admissibility a necessary structural prerequisite for local exponential stability.

Persistence $\not\Rightarrow$ stability

Persistence is **neither necessary nor sufficient** for stability. A persistent orientation can still destabilise the loop; the stability certificate depends on target geometry and edge weights.

SDP weight design

A semidefinite program finds edge weights maximising the exponential convergence rate without altering the sensing graph.

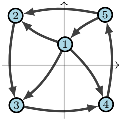

$$\begin{bmatrix}(p_1^*)^\top\\(p_2^*)^\top\\(p_3^*)^\top\\(p_4^*)^\top\\(p_5^*)^\top\end{bmatrix} = \begin{bmatrix}0.0239 & 0.4636\\-1.0826 & 1.1027\\-1.0627 & -0.8972\\0.9372 & -0.8775\\0.9173 & 1.1224\end{bmatrix}$$

Admissible wheel orientation

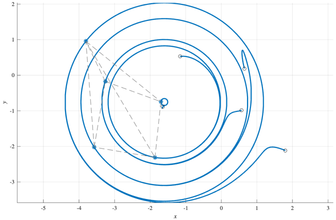

Persistent but unstable

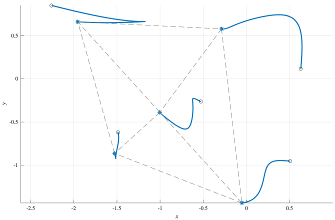

SDP weights: convergent