From: Joe Maxwell from Joe Maxwell
joesystemica@substack.com

To: joseph.maxwell02@proton.me
joseph.maxwell02@proton.me

On: Tuesday, 26 May 2026 at 3:08:45

Minimal Topology of Perturbed Dynamic Systems

Constraint, state space, and the emergence of geometry

JOE MAXWELL
MAY 26

Abstract

Complex systems across many domains often exhibit a similar behavioural pattern: sustained stability under load, progressive loss of recovery capacity, and abrupt transition once a threshold is crossed. Biological regulation, neural systems, engineered infrastructure, financial systems, materials, and cognitive systems can all display this structure, despite differing in their underlying mechanisms [1–4].

This paper proposes a constraint-first description of such behaviour. Instead of beginning with component-level mechanisms, systems are described as trajectories moving through constrained state spaces. Constraints define the reachable regions of the state space, while accumulated perturbation drives trajectories toward boundaries between regimes. When trajectories approach or cross these boundaries, abrupt transitions, hysteresis, and collapse-like behaviours may occur [2,3,5].

A minimal coupling topology is introduced for regulatory systems composed of interacting subsystems. A triadic subsystem model is insufficient when the interaction between subsystems produces an

integrated system state that feeds back into the subsystems themselves. Adding this integrated state yields a four-node fully coupled structure, represented by the complete graph K4, geometrically equivalent to a tetrahedral simplex.

Within this framework, geometry is not treated as a prior assumption. It emerges as the record of constrained interaction: attractor basins, transition surfaces, bottlenecks, and collapse trajectories appear from the relationship between state space, constraint, perturbation, and coupling.

The framework is summarised as:

$(S, C, \Pi, K4) \rightarrow G$

where S is system state space, C is constraint structure, Π is accumulated perturbation, K4 is minimal full coupling topology, and G is the emergent geometry governing system behaviour.

The aim is not to replace existing dynamical systems theory, nor to derive new universal equations of motion. It is to identify a minimal structural grammar for describing how regulatory systems under constraint stabilise, drift, collapse, and reorganise.

1. Systems under pressure

Across domains, complex systems often appear stable until they are not.

A material can carry increasing load before yielding. A queue can absorb demand before backlog becomes unbounded. A physiological system can compensate for stress before dysregulation becomes visible. An AI conversation can remain coherent before drifting into unsupported claims. An infrastructure system can continue operating while internal recovery capacity is quietly being consumed.

These transitions are often described after the fact as failures, collapses, breakdowns, or phase changes. The visible event may appear sudden, but the underlying trajectory is usually not instantaneous. A system can move gradually toward a boundary while continuing to produce acceptable outputs.

This motivates a different starting point.

Instead of asking only:

which component failed?

we can ask:

how did the system's trajectory move through its constrained state space until its prior regime was no longer recoverable?

This is the core move.

The minimal topology paper frames the same shift directly: collapse phenomena should not be understood only through component failure, but through the structural organisation of the system as a whole and the geometry of its possible trajectories.

2. Constraint-first modelling

Let a system be represented by a state vector $x(t)$, where $x(t)$ belongs to a state space S.

The state space S represents the possible configurations the system may occupy.

In standard dynamical form, system evolution can be written as:

$$dx/dt = F(x, t)$$

where F represents the internal dynamics governing motion through the state space.

However, not all states in S are reachable. Systems operate under constraints. These constraints may be physical, energetic, informational, structural, social, biological, or computational.

Let $C(x)$ define a constraint surface.

The constrained region of the state space can be written as:

$$S\_C = \{x \text{ in } S \mid C(x) \leq 0\}$$

This means the system may only occupy states that satisfy the active constraint condition.

Under this framing, behaviour is not determined by internal dynamics alone. It is determined by the interaction between:

1. the current state of the system
2. the internal dynamics

3. the active constraint set
4. the accumulated perturbation acting on the system

The important point is simple:

constraints do not merely limit behaviour from outside; they shape the geometry of possible behaviour.

A system does not move through an empty space.

It moves through a constrained one.

This is consistent with classical dynamical systems thinking, where trajectories, attractors, stability, and bifurcations describe how systems evolve under governing equations and changing conditions [3]. The difference here is emphasis: the framework begins with the geometry of constraint rather than with a specific domain mechanism.

3. Perturbation and accumulated pressure

Most systems do not collapse because of one isolated input.

They collapse because perturbation accumulates, recovery becomes insufficient, and the system approaches a boundary where its previous dynamics no longer hold.

Let $p(t)$ represent instantaneous perturbation.

Accumulated perturbation can be represented as:

$\Pi(t) = \int p(\tau)d\tau$ from 0 to t

In words, $\Pi(t)$ is the pressure history of the system.

This matters because two systems exposed to the same immediate perturbation may respond differently depending on prior accumulated load. A body after weeks of poor sleep does not respond like a rested body. A material with accumulated microstructural damage does not respond like a pristine material. A queue after sustained backlog does not respond like an empty queue. A model conversation after repeated ambiguity does not respond like a clean prompt.

Perturbation history changes state.

As Π increases, the trajectory of the system may move closer to constraint boundaries. Near such boundaries, small additional perturbations can create large behavioural changes. From the outside, this looks like sudden collapse. From the state-space view, it is boundary crossing.

The original formulation describes this clearly: accumulated perturbation gradually pushes trajectories toward constraint boundaries, and when boundaries are crossed, systems transition to new operating regimes.

4. Relation to existing dynamical systems theory

This framework is not intended to replace existing dynamical systems theory.

It sits beside it as a structural grammar.

Classical dynamical systems theory already provides tools for analysing trajectories, attractors, stability, bifurcations, and nonlinear transitions [3]. Catastrophe theory describes how continuous changes in control parameters can produce abrupt qualitative changes in system state [5]. Critical-transition research describes how complex systems may show early-warning signals, hysteresis, and regime shifts before collapse or reorganisation [2].

The contribution here is narrower.

It asks what minimal structural ingredients are required to describe regulatory systems under accumulated perturbation, especially when multiple subsystems interact and produce an integrated state that feeds back into the system.

In this sense, the framework complements:

- attractor analysis
- bifurcation theory
- catastrophe models
- nonlinear dynamics
- feedback-control thinking
- active-inference and free-energy style state-space approaches [6]

The purpose is not to derive a new universal law of motion.

It is to clarify the minimal topology through which constrained regulatory systems can be represented.

That claim ceiling matters.

Without it, the framework becomes too broad.

With it, the framework becomes testable.

5. State-space geometry

Once a system is described in terms of constrained state spaces, geometric structure begins to matter.

A stable operating regime can be represented as an attractor basin: a region of state space in which trajectories tend to remain bounded or return after perturbation.

A transition boundary separates one regime from another.

A bottleneck restricts possible trajectories.

A collapse trajectory describes motion away from a prior basin after recovery fails.

The geometry of the constrained state space therefore determines:

- where stability occurs
- how drift accumulates
- which transitions become possible
- how collapse unfolds
- whether recovery can return to the prior regime or must enter a new one

This does not require geometry to be assumed as a physical substrate.

It means geometry appears as a description of allowed motion.

That is the key shift.

Geometry is not the cause of emergence.

Geometry is the record of constraint.

The source paper makes this statement explicitly: geometry does not cause system behaviour; it records how constraints and subsystem interactions shape the available trajectories.

6. Minimal coupling topology

Regulatory systems are rarely single-variable systems.

They are usually composed of interacting subsystems.

Consider a triadic regulatory system with three interacting components:

A, B, C

These may represent many domain-specific structures. For example:

- sensing, modelling, control
- body, cognition, regulation
- generation, transmission, feedback
- substrate, processor, signal layer
- material, load path, environment

A triadic interaction structure can support nonlinear feedback loops.

However, a triad alone does not explicitly represent the integrated global state produced by subsystem interaction.

If A, B, and C interact continuously, their interaction can produce an emergent integrated system state:

$H = h(A, B, C)$

This integrated state H is not just another independent component. It is the global configuration produced by the interaction of the subsystems.

Once H feeds back into A, B, and C, the system contains four mutually interacting nodes:

A, B, C, H

If each node can influence the others, the interaction graph becomes the complete graph on four vertices:

$K4$

$K4$ has:

- four nodes

- six edges
- full pairwise coupling
- a tetrahedral geometric realisation

This is why K4 appears as a minimal topology.

A system with only two nodes can represent pairwise interaction.

A system with three nodes can represent triadic feedback.

But a system with three subsystems plus an integrated system state requires four nodes.

If that integrated state feeds back into the subsystems, the minimal fully coupled topology is K4.

The original paper states this directly: K4 is the smallest topology capable of modelling interacting subsystems, an emergent integrated state, and full feedback coupling.

7. The structural relationship

The framework can be compressed into one expression:

$(S, C, \Pi, K4) \rightarrow G$

where:

S = system state space
C = constraint structure
Π = accumulated perturbation
K4 = minimal fully coupled subsystem topology
G = emergent geometry of system behaviour

This expression means:

when a system moves through a constrained state space under accumulated perturbation, and when its subsystems are fully coupled through an integrated state, the available behaviours form an emergent geometry.

That geometry includes:

- attractor basins
- transition surfaces

- bottlenecks
- collapse trajectories
- recovery paths
- new stable regimes

The expression should not be read as a physical claim that all systems literally contain tetrahedra.

It is a structural grammar.

It says that once an interacting triad produces an integrated state that feeds back into the triad, a four-node fully coupled topology is the minimal representation of that regulatory structure.

The geometry is therefore not decorative.

It is the minimal map of the coupling.

8. Collapse regimes

Under increasing perturbation, trajectories approach the boundaries of stable basins.

Near these boundaries, small disturbances can produce disproportionate changes. The system may then transition into one of several broad collapse regimes.

The current framework identifies four primary collapse morphologies.

Freeze

Freeze occurs when constraint saturation prevents exploration of alternative trajectories.

The system cannot move freely because too many degrees of freedom are blocked. It becomes rigid, immobile, or functionally paralysed.

Examples may include overloaded cognition, saturated queues, bureaucratic gridlock, motor shutdown, or systems where every available action worsens constraint.

The system does not explode.

It locks.

Nova

Nova occurs when feedback amplifies perturbation.

A disturbance increases instability, which increases the disturbance, which increases instability again. The result is runaway amplification.

Examples include thermal runaway, cascading infrastructure failure, escalating panic, financial spirals, or runaway hallucination in a language model where one unsupported claim becomes context for the next.

The system accelerates away from stability.

Inevitability

Inevitability occurs when recovery pathways become progressively blocked.

The system continues moving, but the available paths increasingly point toward degradation. Unlike Nova, the dominant feature is not explosive runaway but progressive narrowing of viable alternatives.

The collapse becomes hard to prevent because the system's own trajectory removes future recovery options.

Adaptive reorganisation

Adaptive reorganisation occurs when the system leaves a prior basin but stabilises in a new one.

This may appear as collapse from the perspective of the old regime, but as transition from the perspective of the larger state space.

A system may not recover by returning to its previous state.

It may recover by reorganising around a new attractor.

The source paper maps these regimes onto a tetrahedral collapse space, treating them as broad trajectories within the geometry produced by subsystem coupling and constraint dynamics.

This is useful because collapse type determines intervention.

Freeze needs degrees of freedom.

Nova needs feedback suppression.

Inevitability needs broken reinforcement paths.

Adaptive reorganisation needs guided transition.

9. Structural recurrence

Across several modelling layers, a recurring sequence appears:

1 → 3 → 1 → 4

This means:

1. the system is first treated as a unity
2. it is decomposed into three interacting subsystems
3. subsystem interaction produces a new integrated unity
4. that integrated unity creates a four-node coupled topology when fed back into the subsystems

This can be written as:

Unity → triadic decomposition → emergent unity → tetrahedral topology

A compressed symbolic form is:

1Q → 3X → 1G → 4D

where:

Q = undifferentiated system state
3X = three interacting axes or subsystems
G = integrated geometric structure
4D = three structural coordinates plus one axis of evolution

This expression should be interpreted carefully.

It is not a claim about the fundamental dimensionality of physical spacetime.

It is a structural grammar of constrained modelling.

It says that when a system is first treated as a whole, then decomposed into three interacting variables, then reintegrated into a coherent geometry, the resulting dynamics require a fourth axis: evolution through time or regime.

The original paper protects this claim by saying the expression is not intended as a statement about fundamental physical spacetime, but as a symbolic representation of structural progression within constrained dynamical systems.

That restraint is essential.

Without it, the expression overclaims.

With it, it becomes a useful modelling grammar.

10. Illustrative coupled system

To make the structure concrete, consider a simplified coupled system with three interacting variables:

$x(t)$, $y(t)$, $z(t)$

Let the internal dynamics be:

$$dx/dt = -x + a \cdot y$$

$$dy/dt = -y + b \cdot z$$

$$dz/dt = -z + c \cdot x$$

where a, b, and c represent coupling strengths.

When coupling remains within stable ranges, the system may converge toward a stable attractor.

Now introduce accumulated perturbation Π into one channel:

$$dx/dt = -x + a \cdot y + \alpha \cdot \Pi$$

where α represents sensitivity to perturbation.

As Π increases, the trajectory shifts. If the perturbation pushes the system toward the boundary of its attractor basin, small fluctuations may become sufficient to move the system into a new regime.

The specific behaviour depends on parameter values.

The structural point is more general:

gradual perturbation accumulation can produce abrupt regime transition when a trajectory reaches a geometric boundary in constrained state space.

This illustrative system appears in the original appendix as a minimal example of how constrained dynamical systems can produce regime transitions through coupling, perturbation accumulation, and basin-boundary crossing.

11. Cross-domain interpretation

If the framework is useful, similar structural patterns should appear across domains even when the mechanisms differ.

For example:

In materials, drift may correspond to defect accumulation, microcrack growth, plastic strain, or thermal stress accumulation.

In physiology, drift may correspond to autonomic instability, recovery debt, immune activation, sleep disruption, or metabolic strain.

In cognition, drift may correspond to unresolved load, sensory overload, transition cost, masking debt, or loss of regulatory capacity.

In infrastructure, drift may correspond to backlog, latency, overload propagation, or declining throughput.

In AI reasoning, drift may correspond to accumulating unsupported assumptions, prompt ambiguity, context contamination, or narrative takeover.

The point is not that these systems are identical.

They are not.

The point is that each can be described as a trajectory through constrained state space under accumulated perturbation.

That allows comparison without reduction.

This is consistent with general systems theory's original ambition to study structural features that recur across domains while preserving the need for domain-specific mechanisms [1].

12. Testing and falsification

A framework this broad must be able to fail.

Minimal testable predictions include:

1. Pre-collapse indicators should appear before transition in

measurable systems.

2. Recovery time should increase as systems approach boundary regions.
3. Variance and sensitivity should increase under accumulated perturbation.
4. Collapse morphology should vary by topology and dominant constraint.
5. Recovery path should depend on collapse morphology.

The experimental protocol document formalises this as a cross-domain testing strategy: define state variables, apply controlled perturbation, measure recovery time, variance and sensitivity, increase load, detect transition, observe collapse morphology, and compare to prediction.

The framework weakens if:

- collapse occurs without precursor signals under proper measurement
- stability is independent of drift and recovery rates
- no identifiable basins or boundaries exist
- collapse morphology does not vary structurally
- recovery path is unrelated to collapse type

This matters because the framework must remain a modelling scaffold, not an unfalsifiable vocabulary.

If no measurable pre-collapse signals appear, the framework must be restricted.

If collapse morphology does not improve recovery reasoning, the classification must be revised.

If K4 does not add explanatory power over simpler coupling structures, it should be demoted.

That is the correct standard.

13. Scope and claim ceiling

This paper does not claim that all systems are the same.

It does not claim that K4 is metaphysically fundamental.

It does not claim that tetrahedral geometry literally governs all collapse.

It does not replace domain-specific mechanisms, experiments, or simulations.

It proposes a minimal structural grammar for systems where:

- states are distinguishable
- transitions are constrained
- perturbation accumulates
- subsystems interact
- integrated state feeds back into subsystem behaviour
- regimes exist
- collapse and recovery are trajectory-dependent

Within that scope, the framework may be useful.

Outside it, it should be revised or discarded.

14. Conclusion

Complex systems often appear stable until they abruptly transition into new regimes. Component-level analysis can explain much of system behaviour, but may miss the structural conditions under which interacting subsystems lose stability under accumulated perturbation.

A constraint-first description begins with state space, active constraints, and perturbation history. It treats behaviour as trajectory through a constrained possibility space.

When regulatory systems contain interacting subsystems whose interaction produces an integrated state, the minimal fully coupled topology becomes K4: three subsystem nodes plus the integrated state, with full feedback coupling.

This gives the compact structural expression:

$(S, C, \Pi, K4) \rightarrow G$

Geometry, in this formulation, is not imposed in advance.

It emerges as the record of constrained interaction.

The value of the framework depends on whether it improves prediction, interpretation, and intervention across domains. It should help identify drift before collapse, distinguish collapse morphologies, and match recovery paths to the structure of failure.

If it does that, it has earned use.

If it does not, it should be constrained, revised, or abandoned.

That is sufficient.

References

[1] von Bertalanffy, L. (1968). *General System Theory*. George Braziller.

[2] Scheffer, M. (2009). *Critical Transitions in Nature and Society*. Princeton University Press.

[3] Strogatz, S. H. (2015). *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press.

[4] Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall.

[5] Thom, R. (1975). *Structural Stability and Morphogenesis*. W. A. Benjamin.

[6] Friston, K. (2010). The free-energy principle: a unified brain theory. *Nature Reviews Neuroscience*, 11, 127–138.

[7] Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27, 379–423.