Distributed Intelligence in Motion: Self-Organizing Cognitive Swarms

A new architecture for autonomous, multi-target missions in contested environments

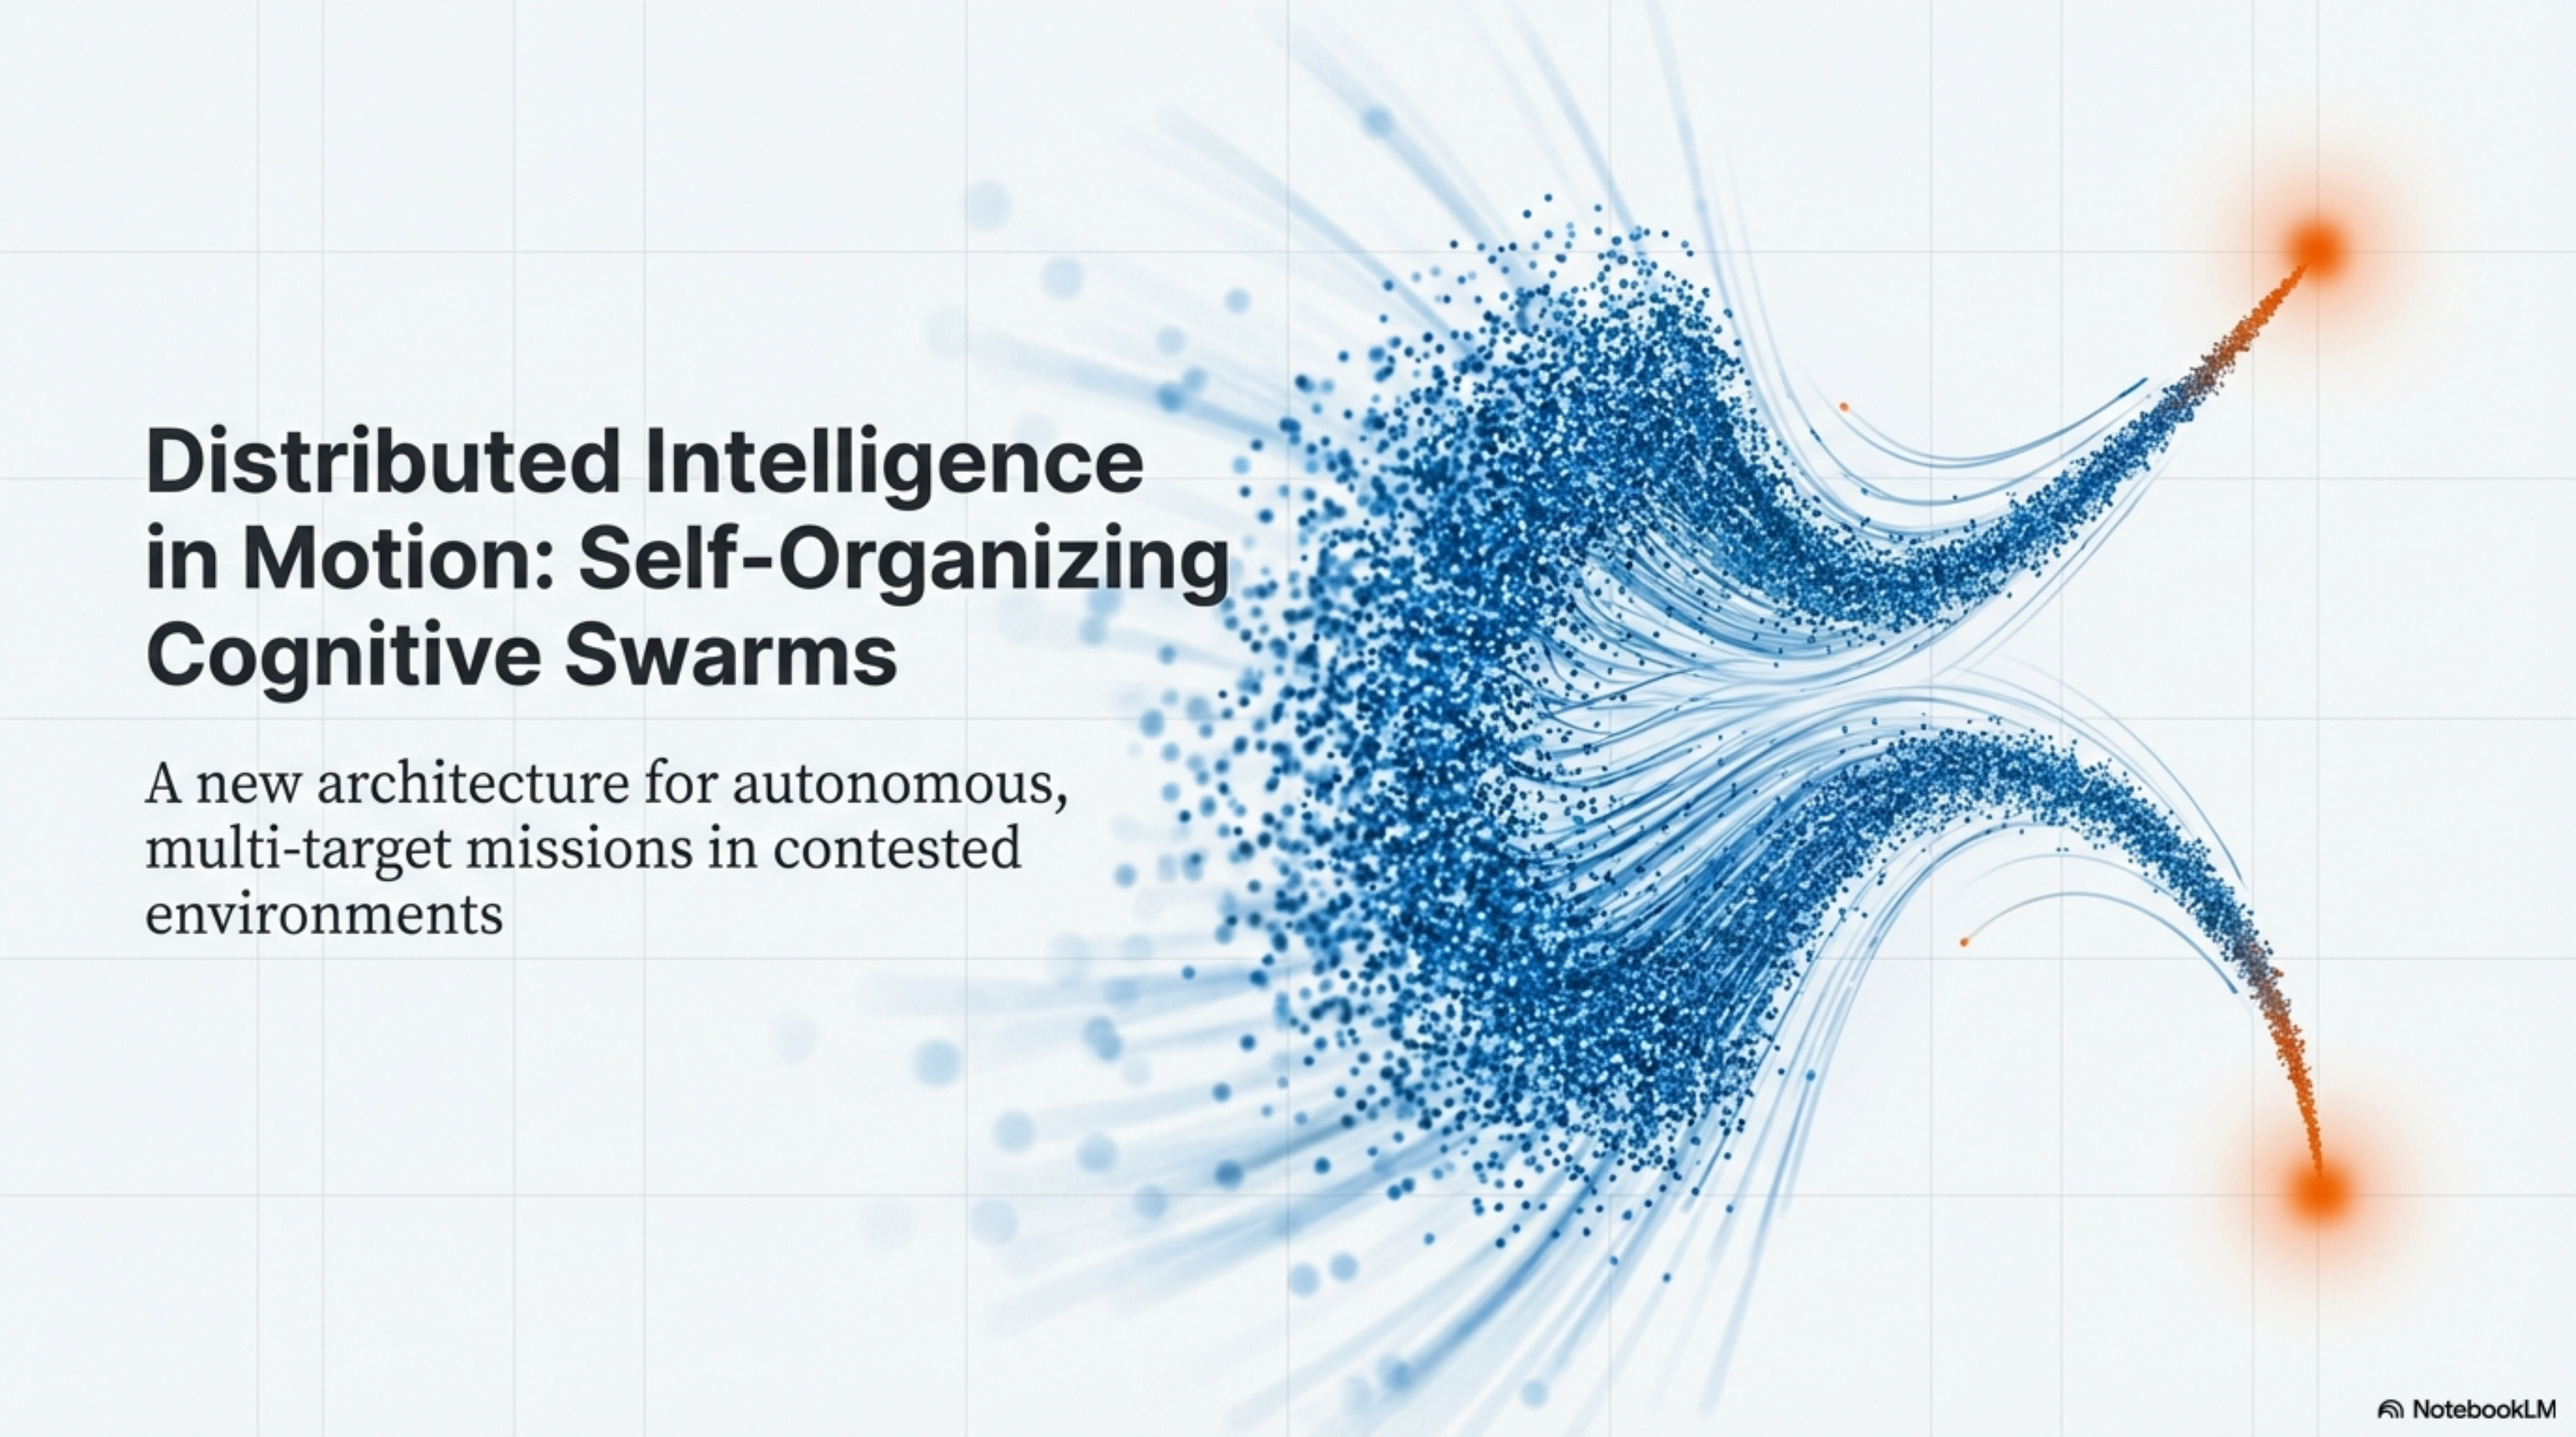

The Demand for True Autonomy has Reached a Tipping Point.

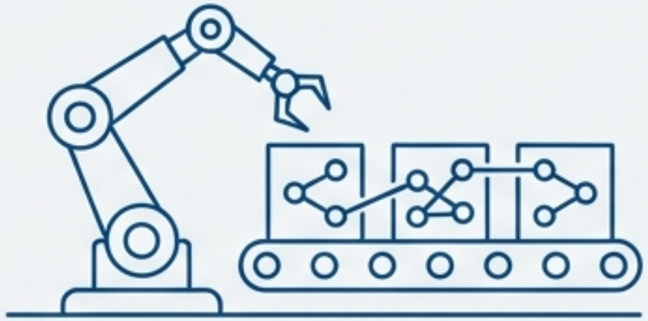

Civil Sector - The Rise of Industry 5.0

- Focus on resilient, cyber-physical-social systems.
- Applications: Large-scale autonomous coordination in logistics, manufacturing, and infrastructure monitoring.
- Core Need: Scalable, robust systems that operate without a central point of failure.

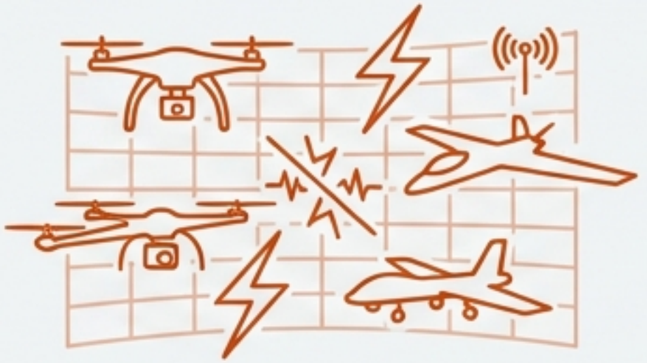

Defense Sector - Lessons from the Modern Battlefield

- Contemporary conflicts reveal the decisive role of autonomous systems under heavy electronic warfare.
- Requires operation where classical remote control is degraded or fully suppressed.
- Core Need: Decentralized intelligence, self-organization, and autonomous mission execution.

"The biggest battlefield innovation in a generation."

— U.S. Department of Defense, July 2025 Memorandum.

Traditional Approaches Break Under Pressure

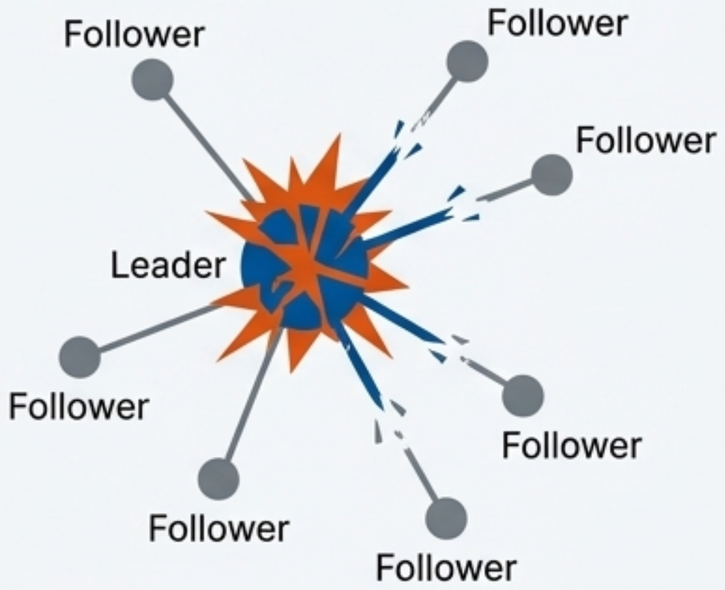

Leader-Follower Architectures

Introduce fragility. The loss of a leader or near-leader node can collapse the entire formation. Creates a single point of failure.

Consensus-Based & Potential-Field Methods

Suffer from slow convergence and oscillations at scale. Vulnerable to perturbations or jamming in dynamic environments.

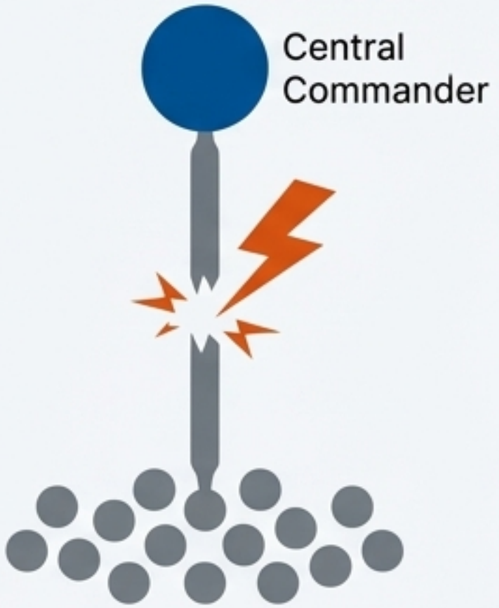

Centralized Control

Fundamentally non-viable in communication-denied or contested environments where a constant link to a central commander cannot be assumed.

How Can a Swarm Think for Itself to Execute Complex Missions?

- How does a large swarm autonomously divide itself to address multiple, spatially distributed targets?
- How does it allocate sufficient resources to each target without a central planner?

- How does it maintain cohesion and mission focus while sustaining heavy losses from adversarial engagement?
- How does it dynamically reconfigure and re-task itself as the mission evolves?

A New Model Where Intelligence Emerges from Structure.

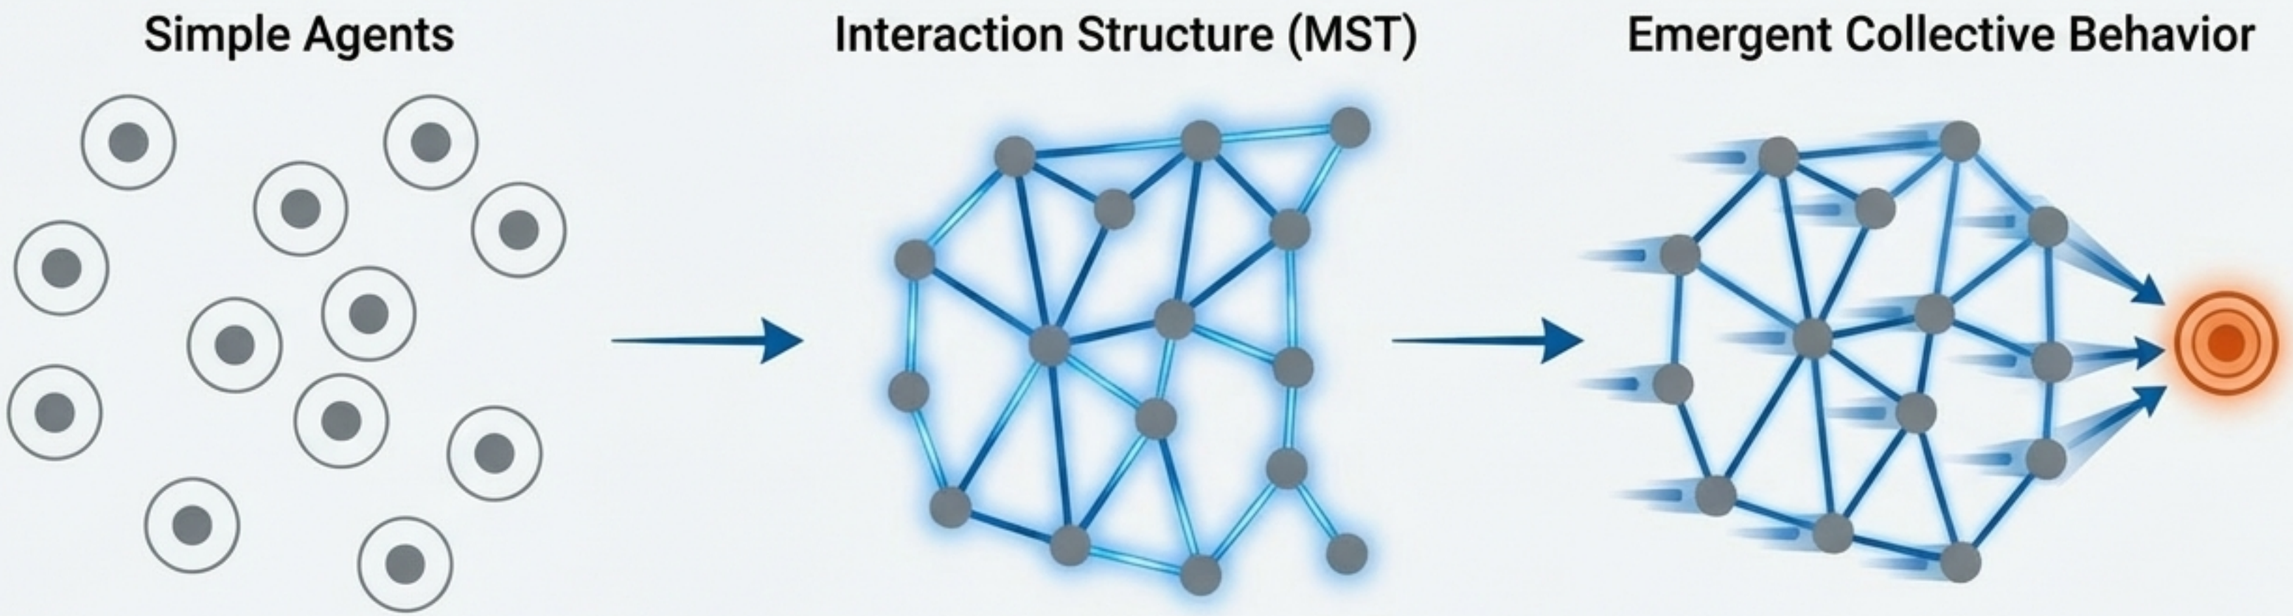

Introducing the Self-Organizing Cognitive Swarm. It operates without a centralized commander. Collective behavior emerges from simple, locally-informed agents. “Cognition” is not located in any single agent but is an emergent property of the collective’s ability to self-organize, adapt, and reconfigure its topology in response to the mission.

> **“...distributed cognitive behavior...driven by local information, collective memory encoded in topology, and continuous reconfiguration of goals.”**

The MST: A Lean and Resilient Nervous System

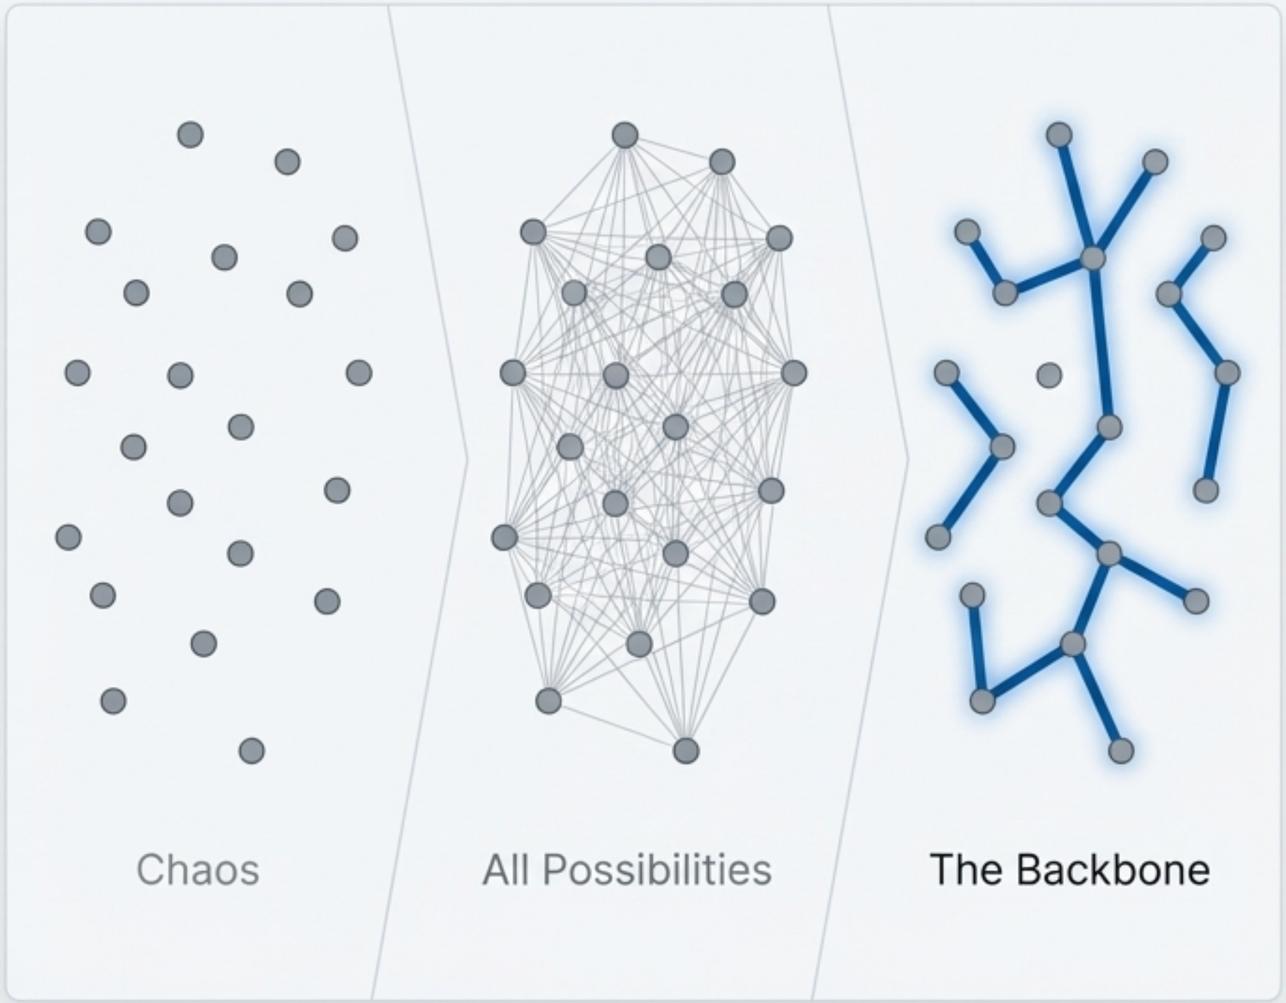

Core Organizational Principle:
The Minimum Spanning Tree (MST) serves as the communication and coordination backbone for each sub-swarm.

Minimal Communication Cost:
The MST ensures connectivity with the minimum possible total edge length, crucial for efficiency in bandwidth-limited environments.

Clear Hierarchy, No Loops: Every agent has a unique parent, ensuring an unambiguous, directed flow of influence from the leader to the entire sub-swarm.

Deterministic Repair: Unlike probabilistic or consensus-based models, the tree structure provides a predictable and interpretable way to recover from agent losses.

Step 1: Divide & Conquer with Balanced Clustering.

A balanced k-means-like clustering mechanism assigns agents to sub-swarms.

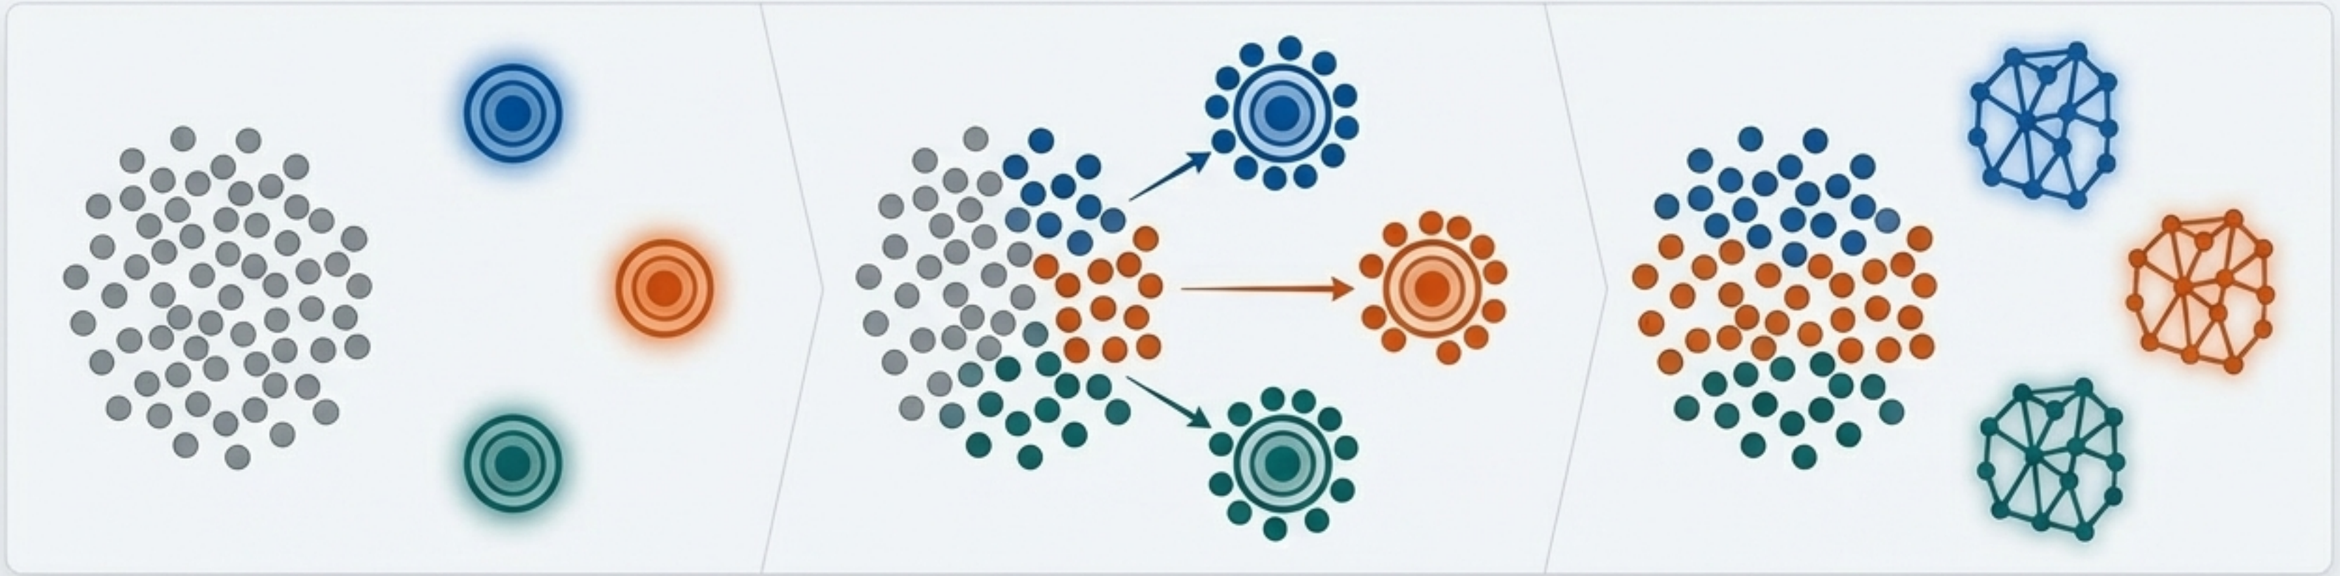

Initial State

Clustering

Sub-Swarm Formation

- **Target-Anchored:** The known coordinates of each target serve as the initial cluster centers.
- **Proportional Allocation:** The algorithm assigns agents to the nearest target-center while a balancing constraint ensures each sub-swarm has a proportionate number of agents.
- **Result:** The swarm autonomously parallelizes the mission, dedicating a properly-sized force to each objective from the outset.

Step 2: Coordinated Motion via Leadership Chains.

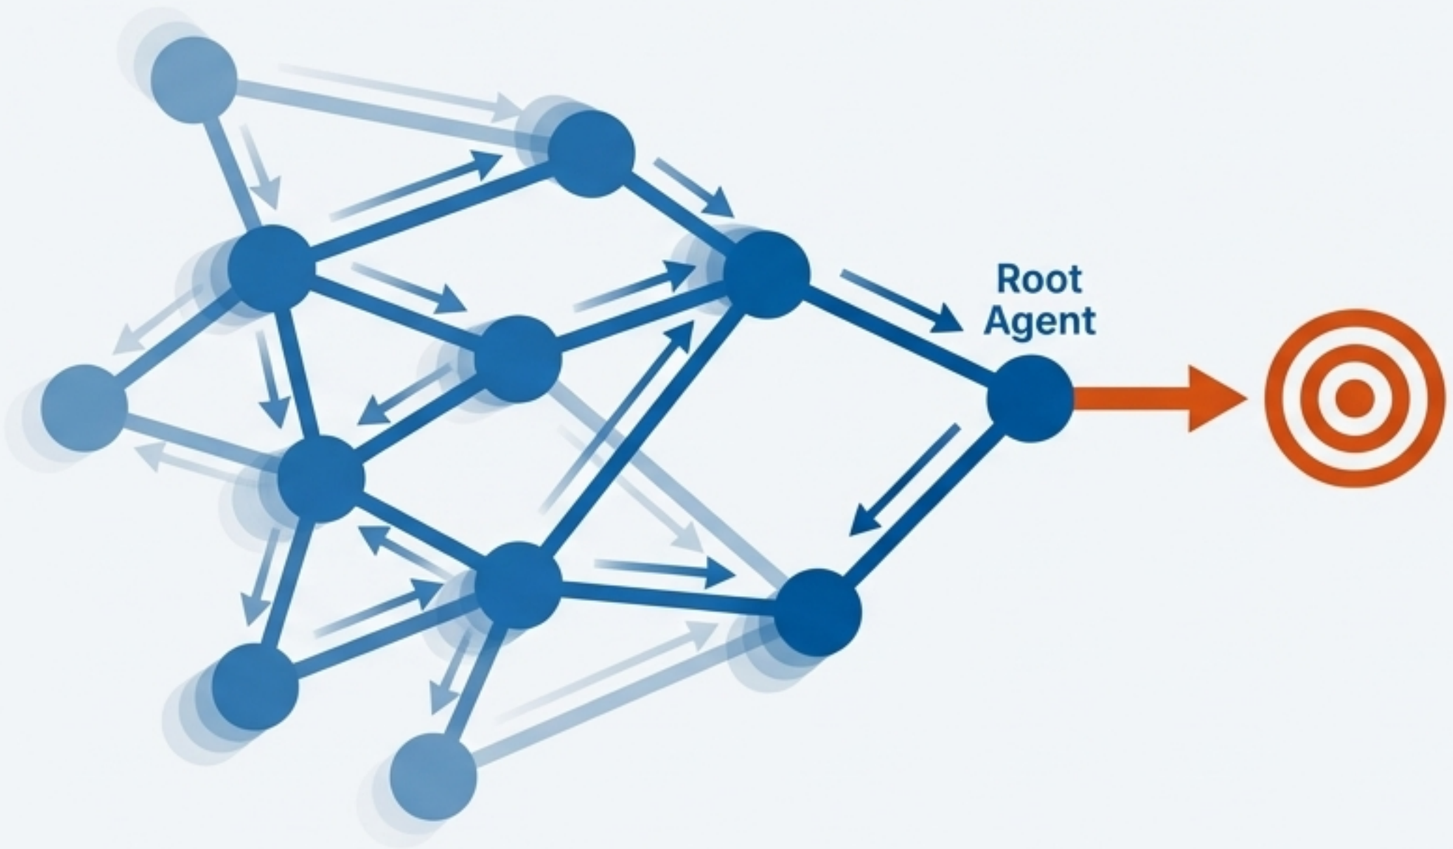

The Root Agent

The leader of each sub-swarm is the agent closest to the target. It is the *only* agent that "sees" the target and moves directly toward it.

$$x_r(t+1) = x_r(t) + \alpha \frac{y_j(t) - x_r(t)}{\|y_j(t) - x_r(t)\|}$$

All Other Agents

Every other agent in the sub-swarm follows a single, simple rule: move toward your designated parent in the MST.

$$x_i(t+1) = x_i(t) + \alpha \frac{x_{L_i}(t)(t) - x_i(t)}{\|x_{L_i}(t)(t) - x_i(t)\|}$$

Emergent Effect

The entire sub-swarm is pulled toward its objective like a deformable chain, with influence propagating deterministically down the MST hierarchy.

Step 3: Strength Through Sacrifice

Attrition Model: “Destructive Cooperation”

Agent's Function

An agent’s primary mission function is to collide with its designated target.

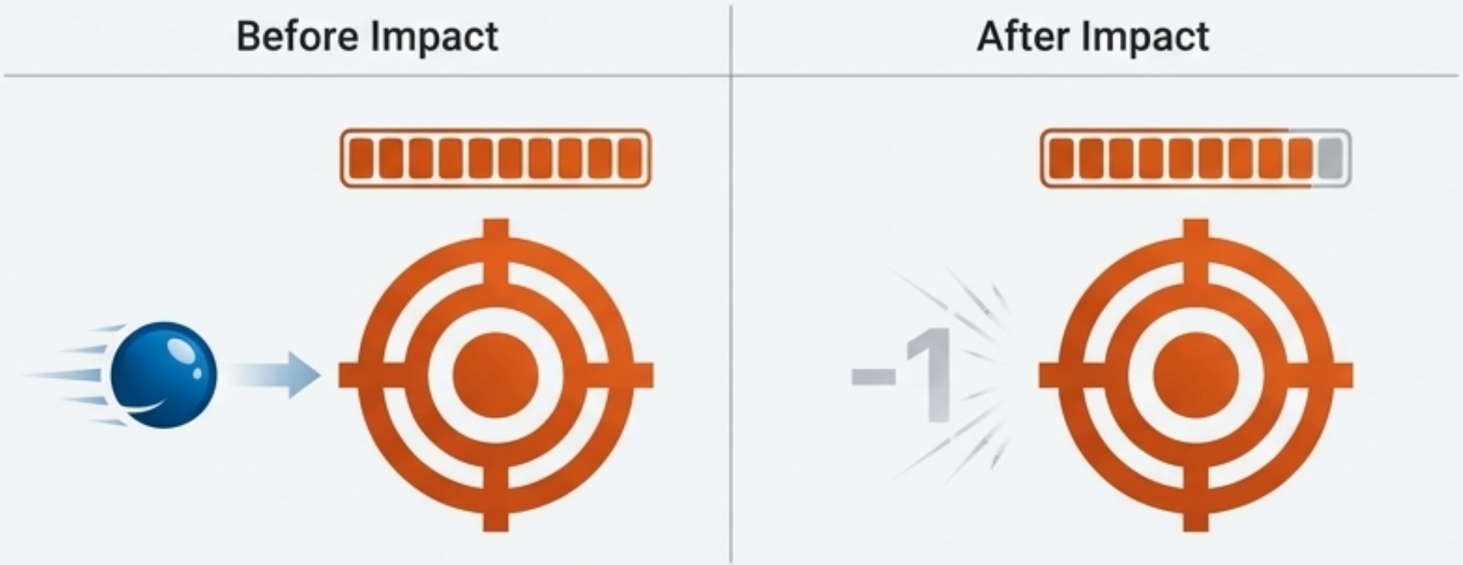

Attrition Dynamics

When an agent collides with a target (distance < δ):

1. The agent is destroyed.
2. The target loses one "life" (durability L_0).
3. The target’s speed degrades proportionally to its remaining lives ($\alpha_{target}(t) = \alpha_T * lives(t) / L_0$).

The Strategic Trade-off

The swarm’s success depends on balancing the rate of sacrifice against the rate of target degradation.

Step 4: When a Link Breaks, the Chain Remakes Itself.

Before

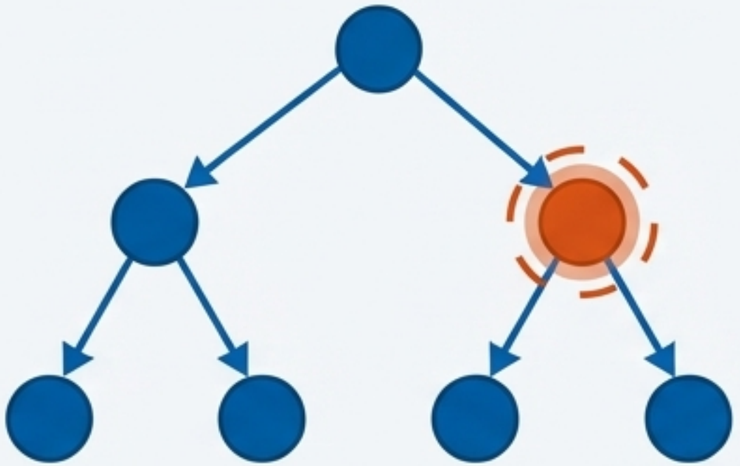

After

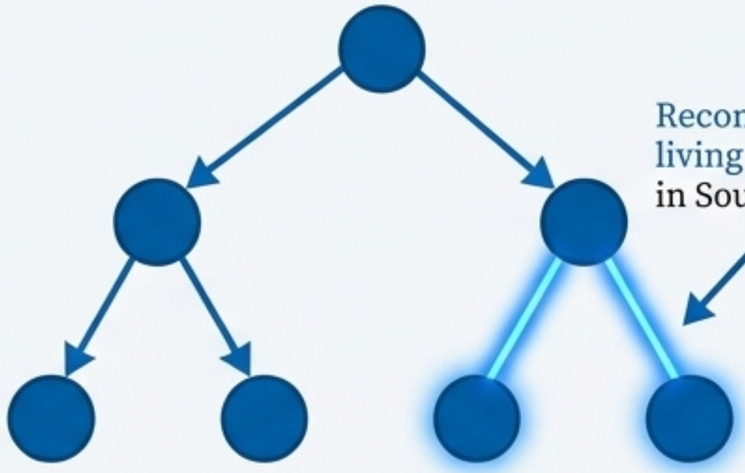

Self-Repair and Reorganization Mechanisms

Agent Loss: If an agent's parent is destroyed, it reconnects to the nearest living ancestor in its parent chain.

Leader Loss: If a root agent is destroyed, one of its surviving children is deterministically promoted. If none exist, the closest surviving agent becomes the new root.

Target Destroyed: The victorious sub-swarm doesn't stop. It merges with the nearest *active* sub-swarm, reinforcing the attack on a remaining target.

Experimental Observations: Emergent Intelligence in Action.

Four Key Principles of Cognitive Swarm Behavior

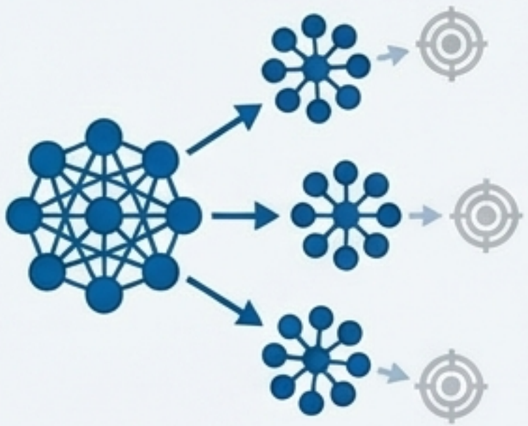

1. Autonomous Parallelization

Mission-oriented sub-swarms reliably form and engage targets without any top-down commands or explicit task allocation.

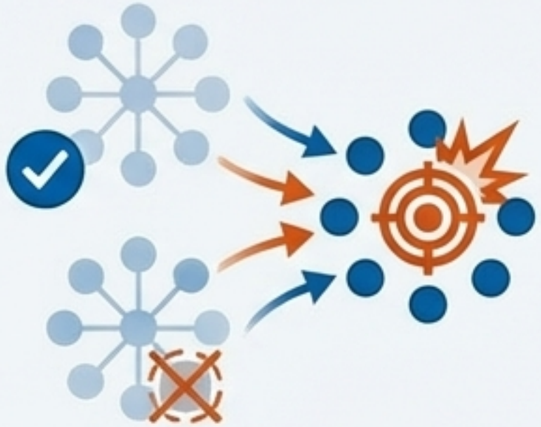

2. Dynamic Load Balancing

After a target is destroyed, victorious sub-swarms autonomously merge and redistribute to reinforce attacks on remaining threats, preventing wasted resources.

3. Graceful Degradation

Swarms maintain cohesion and mission focus even under massive attrition (>50% losses), thanks to deterministic repair mechanisms.

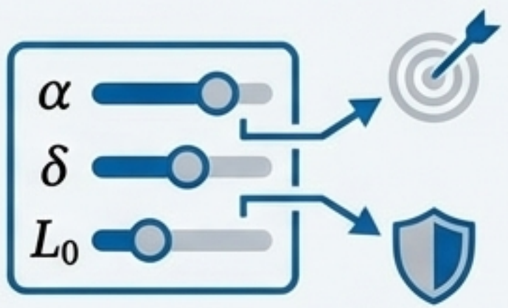

4. Tunable Strategy

The critical trade-off between mission aggression and swarm survivability is directly controllable via a small, interpretable set of hyperparameters.

A New Position on the Spectrum of Autonomy

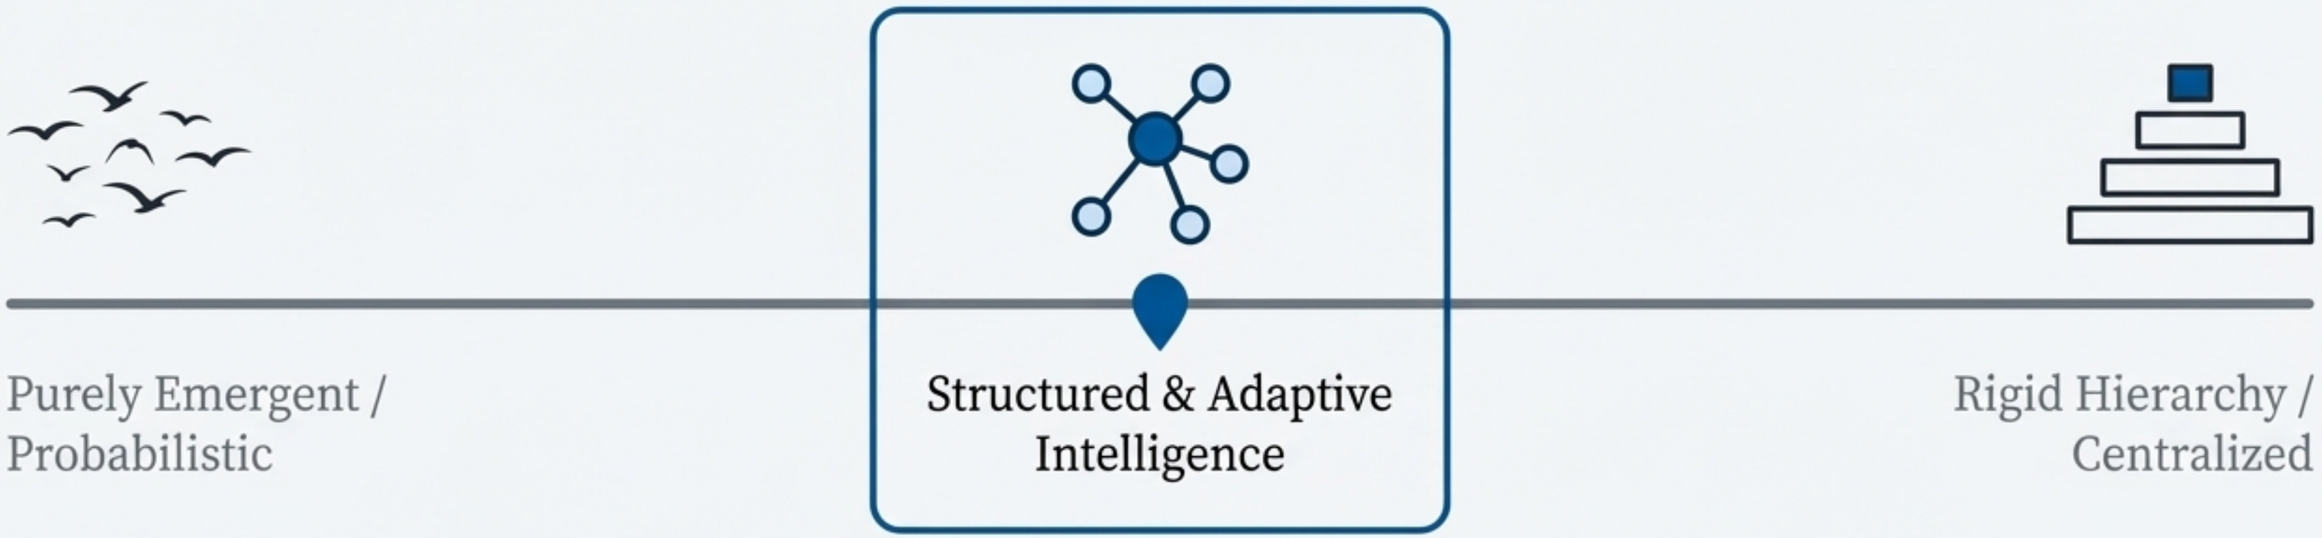

vs. Classical Pursuit-Evasion

Shifts the focus from *optimal agent trajectories* to *organizational survivability* under sustained, sacrificial engagement.

vs. Emergent Flocking (e.g., Reynolds/Vicsek models)

Provides *deterministic recovery* and mission-specific roles via the MST structure, rather than relying solely on probabilistic cohesion. It combines distributed autonomy with explicit resilience rules.

Intelligence is an Organizational Property

Core Contributions

- **A Structurally-Grounded Architecture**
  Resilience and coordination emerge from the MST topology, not from the cognitive complexity of individual agents.
- **Dual Attrition Modeling**
  A novel model of "destructive cooperation" that captures the realistic trade-offs of adversarial engagement.
- **A Reproducible Framework**
  An open-source simulation environment to allow other researchers to explore, stress-test, and extend the proposed mechanisms.

Central Insight

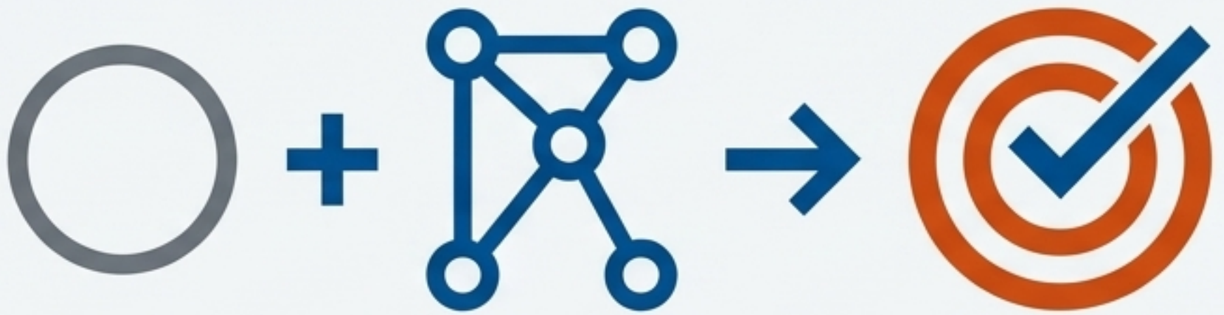

Simple Agent Node + **Resilient Network Structure (MST)** → **Mission Success (Engineered Behavior)**

Complex, mission-oriented behavior can be engineered through simple agents governed by robust, interpretable, and resilient organizational rules.

The Road Ahead: From Simulation to Reality

Current Model: Abstract Simulation

- **Implementation:** Centralized simulation of decentralized concepts for clarity and reproducibility.
- **Physics:** Abstracted sensing and motion models; no energy constraints, communication latency, or aerodynamics.
- **Adversaries:** Targets use simple, reactive avoidance, not strategic adaptation.

Next Steps: Richer Dynamics

- **Modeling:** Incorporate heterogeneous agents, energy constraints, and adaptive adversarial AI.
- **Cognition:** Couple the architecture with learning mechanisms for adaptive hyperparameters or sub-swarm policy learning.

Future Vision: Real-World Deployment

- **Applications:** Apply the framework to non-destructive civil scenarios like disaster response, exploration, and industrial coordination.

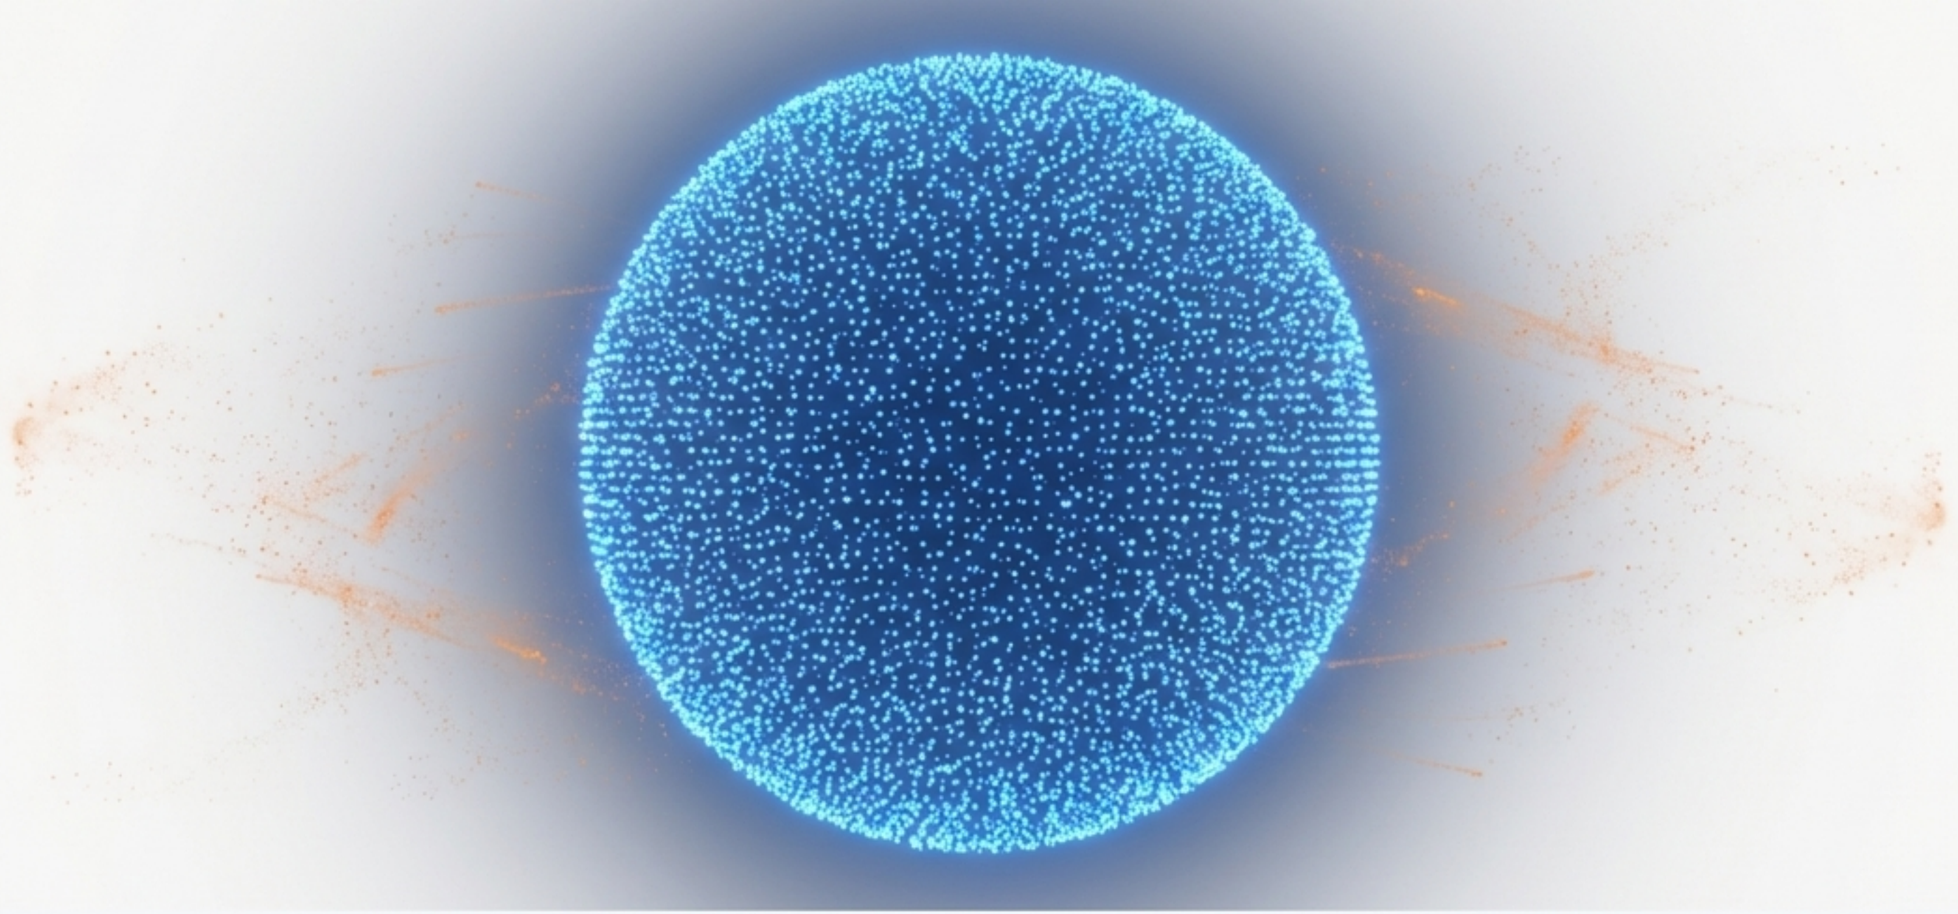

Engineering the Unbreakable Swarm

This work offers a principled pathway toward scalable, resilient, and interpretable multi-agent systems capable of operating where others cannot. It demonstrates that cognitive swarm behavior can be engineered through structure, not complexity.

Read the Full Paper

Explore the Simulation Yourself

https://ai.it.jyu.fi/experiments/swarms/