DARTMOUTH ENGINEERING

Collaboration and Competition in Multi-Agent Systems: The Consequences of Coalitions and Model Mismatch

Bryce L. Ferguson

April 21st, 2026

Game On!

Tragedy of the Commons & Distributed Control

The Grand Banks Fishing Fiasco:

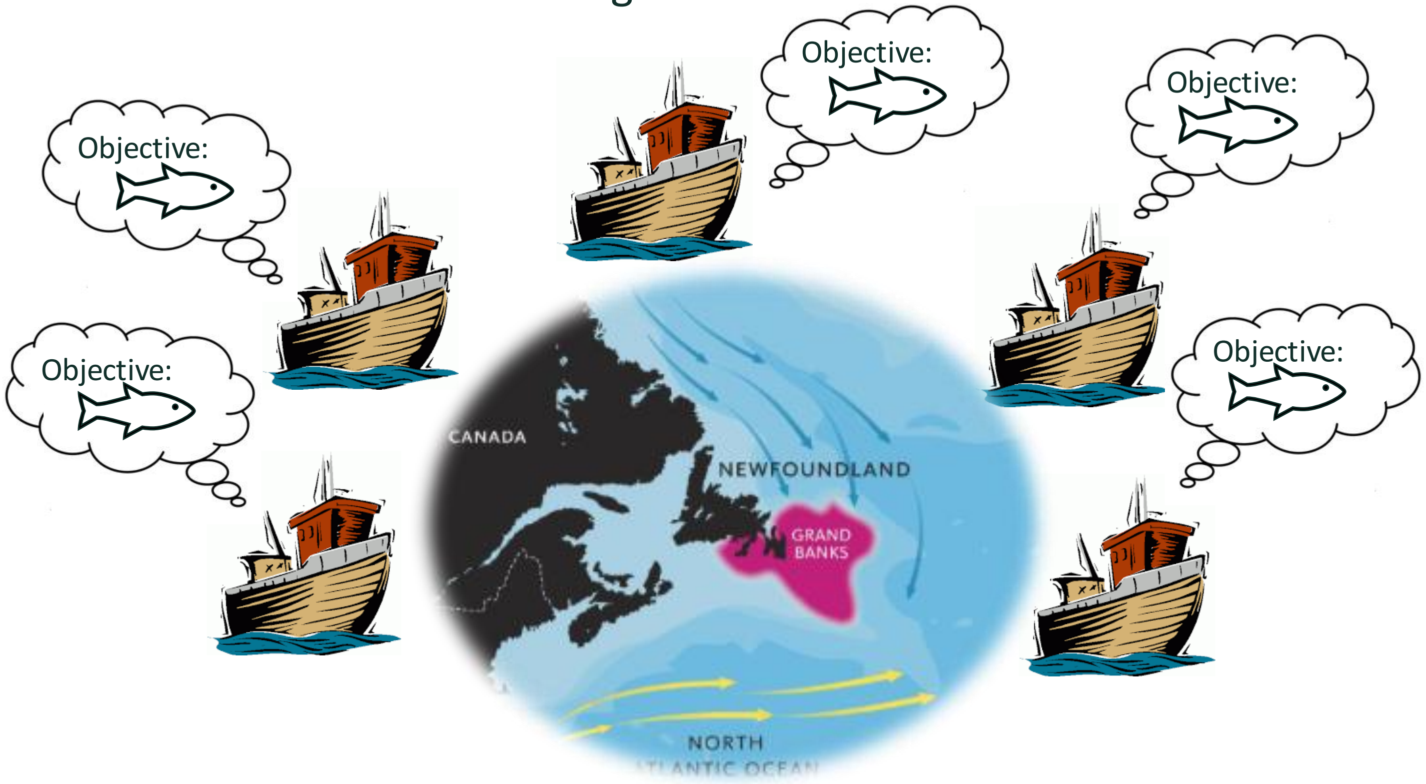

The Collapse of the Grand Banks Cod Fishery

The Grand Banks of Newfoundland are a series of underwater plateaus found off the northeast coast of Canada. For centuries the Grand Banks produced a seemingly endless supply of cod and other commercially valuable fish. However, overfishing and mismanagement of stocks in the second half of the twentieth century saw the number of cod dwindle and then ultimately collapse, causing economic turmoil and significant social issues for the people of Newfoundland who had come to rely on cod fishing as a major part of their economy. As Charles Clover writes in *The End of the Line*:

"The Grand Banks is the textbook case of failure in fisheries science. An army of scientists in one of the world's richest and most advanced nations managed to destroy one of the richest fisheries in the world while convincing themselves for a decade that they were doing no such thing. The Newfoundland cod collapse was the nightmare that shook the world out of its complacent assumption that the sea's resources were renewable..

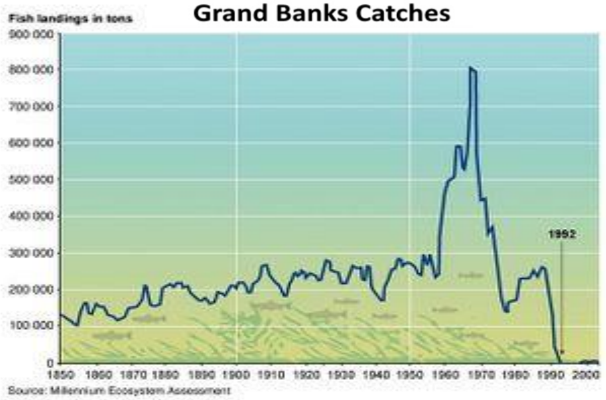

Local Decision Making ⟹ **Sub-optimal Global Behavior**

(catch many fish) (fish population disappears)

The Need for Coordination

Fleet Robotics

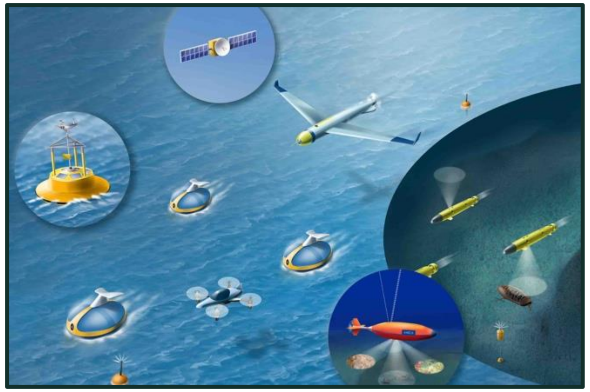

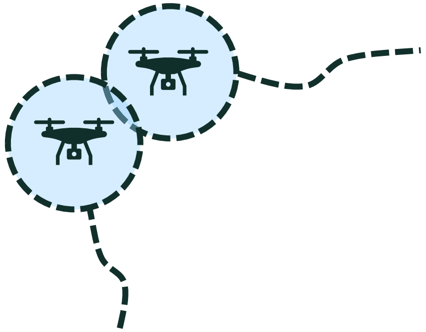

Coordinate…
Local algorithms for many robots

Mass Transportation

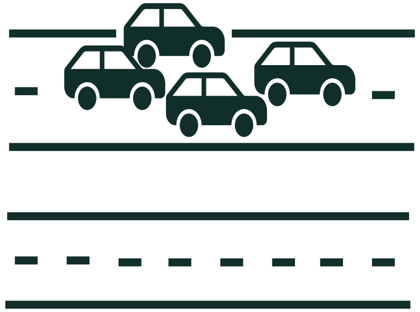

Coordinate…
Decision-making of human users

Energy Consumption

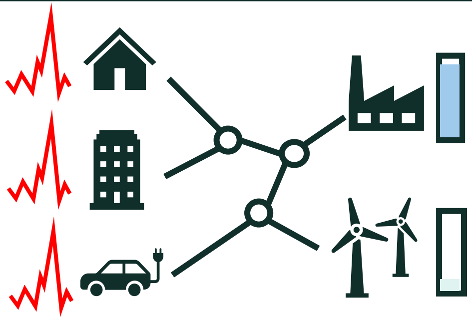

Coordinate…
Generation, demand, and storage of power

Prevalent phenomenon in multi-agent systems

Local Decision Making → **Sub-optimal Global Behavior**

Central Goal

Discover control approaches that are ***locally implementable*** while still offering ***good collective behavior***.

Part I

Agents that can communicate and collaborate with one another

Part II

Noncooperative or competitive agents’ control problem

I. Inter-agent Collaboration

Engineered Multi-Agent Systems

Drone Fleets

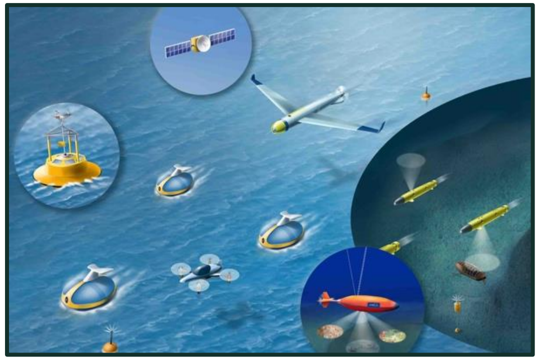

Manufacturing Robotics

Internet of Things

Large, interconnected systems of many automated devices and processes

Quantifying the effects of Collaborative Decision-Making

[CDC23, TAC*, CDC24]

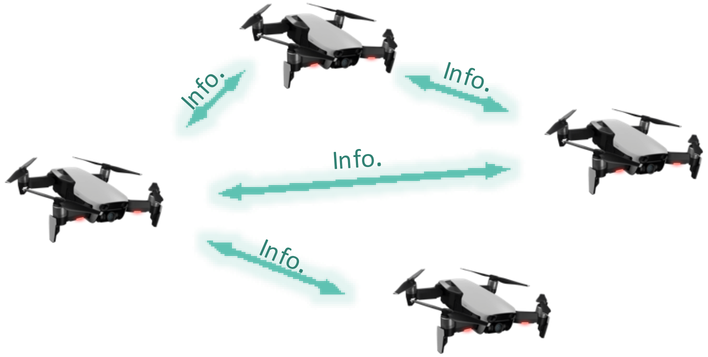

- Enable agents to ***share information and control decisions*** between one another
- Exploit communication channels to elicit greater coordination

Distributed Control
(local decisions)

Design new system paradigms
between centralized and distributed

Centralized Control
(Coordinated decisions)

Modeling Collaborative Solutions

Design Local Control Law

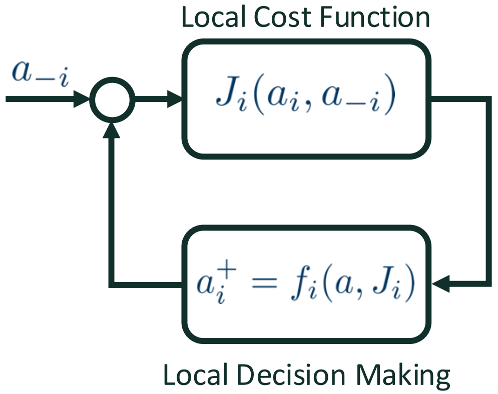

Converge to Nash Equilibria

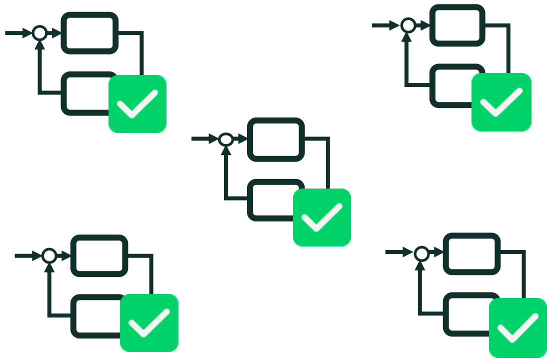

Quality of Nash Equilibria

$$\mathtt{Eff} = \frac{W(a^{\mathrm{NE}})}{W(a^{\mathrm{opt}})}$$

Modeling Collaborative Solutions

Design Local Control Law

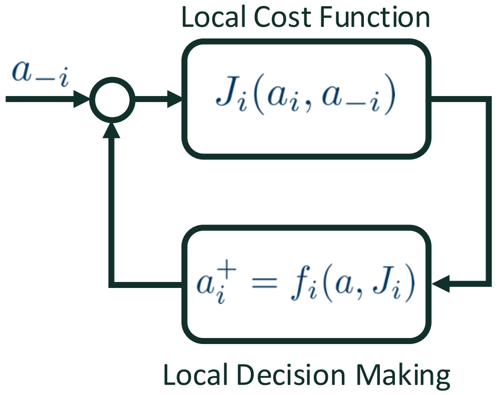

Converge to Nash Equilibria

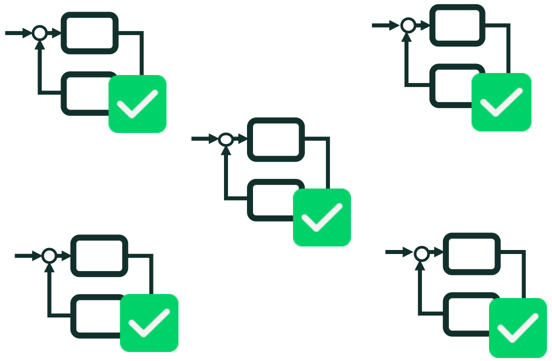

Quality of Nash Equilibria

$$\mathtt{Eff} = \frac{W(a^{\mathrm{NE}})}{W(a^{\mathrm{opt}})}$$

Performance can be ***sub-optimal***

Modeling Collaborative Solutions

Design ***Group*** Control Law

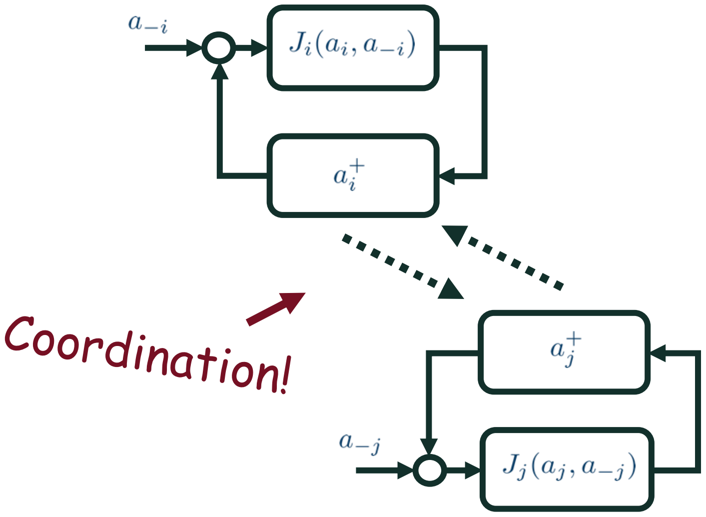

Inter-agent communication
and coordination

Converge to Nash Equilibria

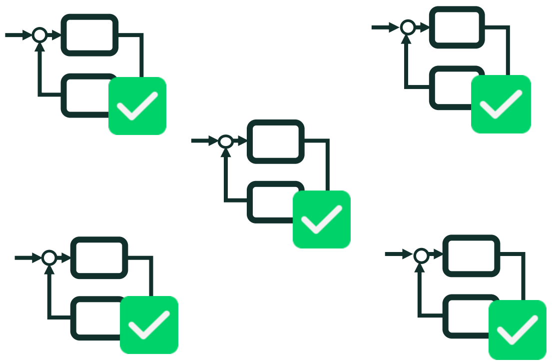

Quality of Nash Equilibria

$$\texttt{Eff} = \frac{W(a^{\mathrm{NE}})}{W(a^{\mathrm{opt}})}$$

Performance can be
sub-optimal

Modeling Collaborative Solutions

Design ***Group*** Control Law

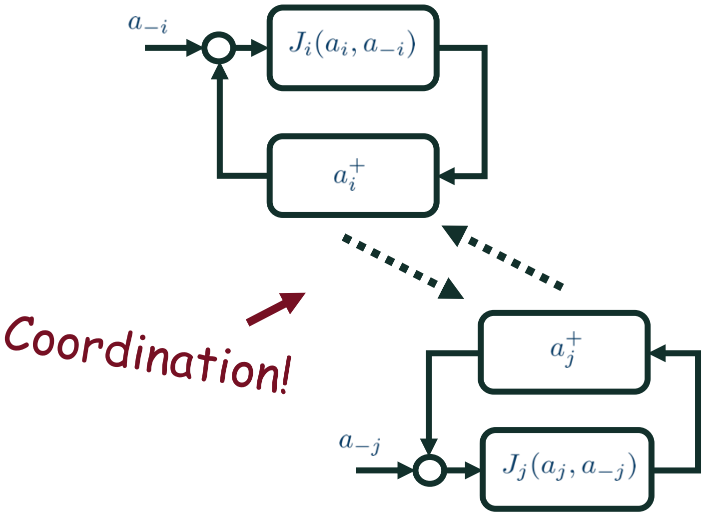

Inter-agent communication and coordination

Converge to ***Stronger*** Equilibria

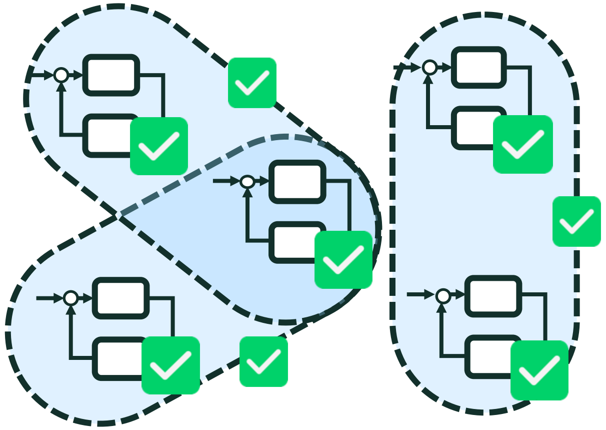

New, stronger equilibria notions

Quality of Nash Equilibria

$$\texttt{Eff} = \frac{W(a^{\mathrm{NE}})}{W(a^{\mathrm{opt}})}$$

Performance can be ***sub-optimal***

Modeling Collaborative Solutions

Design ***Group*** Control Law

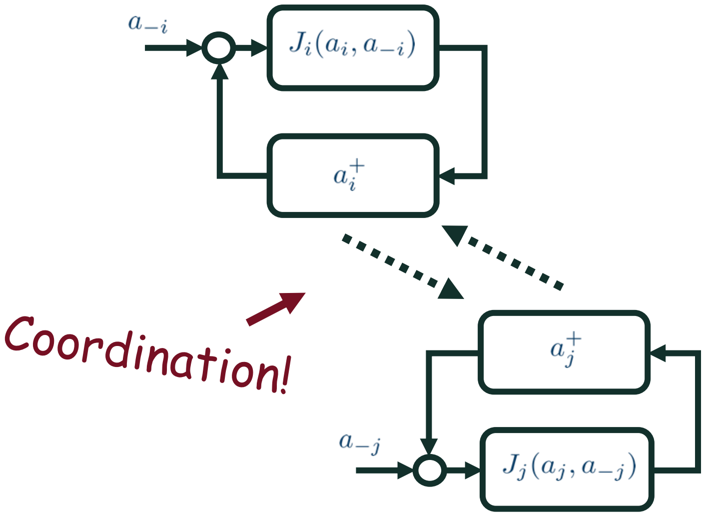

Inter-agent communication and coordination

Converge to ***Stronger*** Equilibria

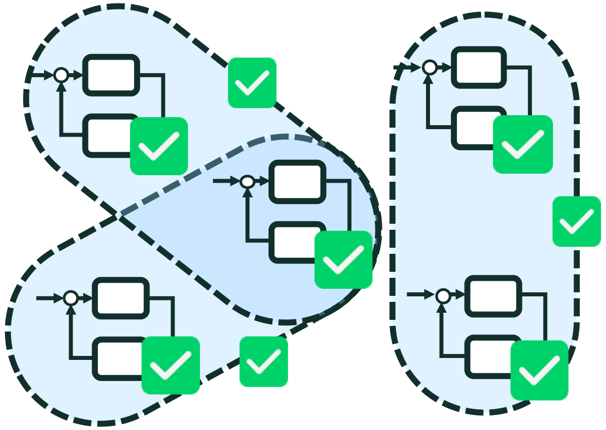

New, stronger equilibria notions

Quality of ***Strong*** Equilibria

$$\frac{W(a^{\text{strong}})}{W(a^{\text{opt}})} \geq \frac{W(a^{\text{NE}})}{W(a^{\text{opt}})}$$

Improve system level performance

Modeling Collaborative Solutions

Design ***Group*** Control Law

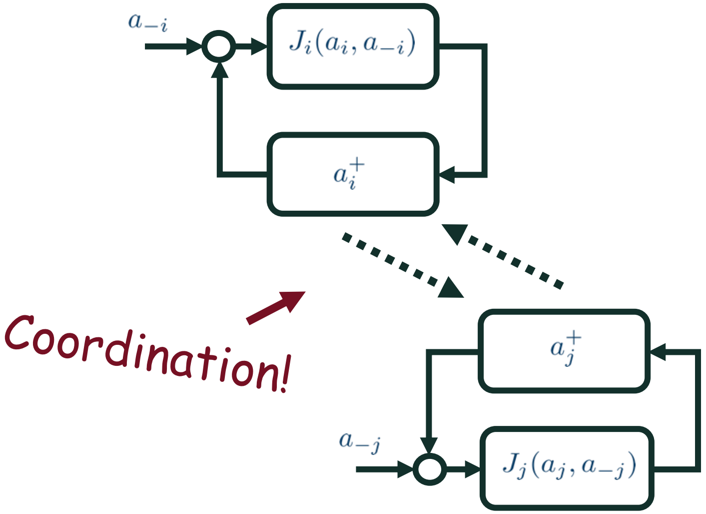

Inter-agent communication and coordination

Converge to ***Stronger*** Equilibria

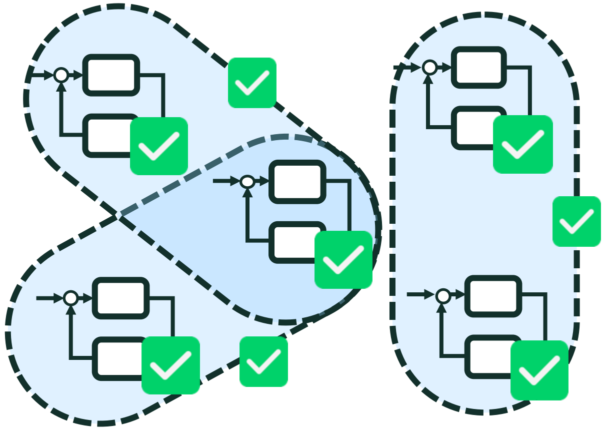

New, stronger equilibria notions

Quality of ***Strong*** Equilibria

$$\frac{W(a^{\mathrm{strong}})}{W(a^{\mathrm{opt}})} \geq \frac{W(a^{\mathrm{NE}})}{W(a^{\mathrm{opt}})}$$

Improve system level performance

With emerging communication technologies, can we go ***beyond local decision-making*** to coordinate better equilibria?

Collaborative Decisions and Emergent Equilibria

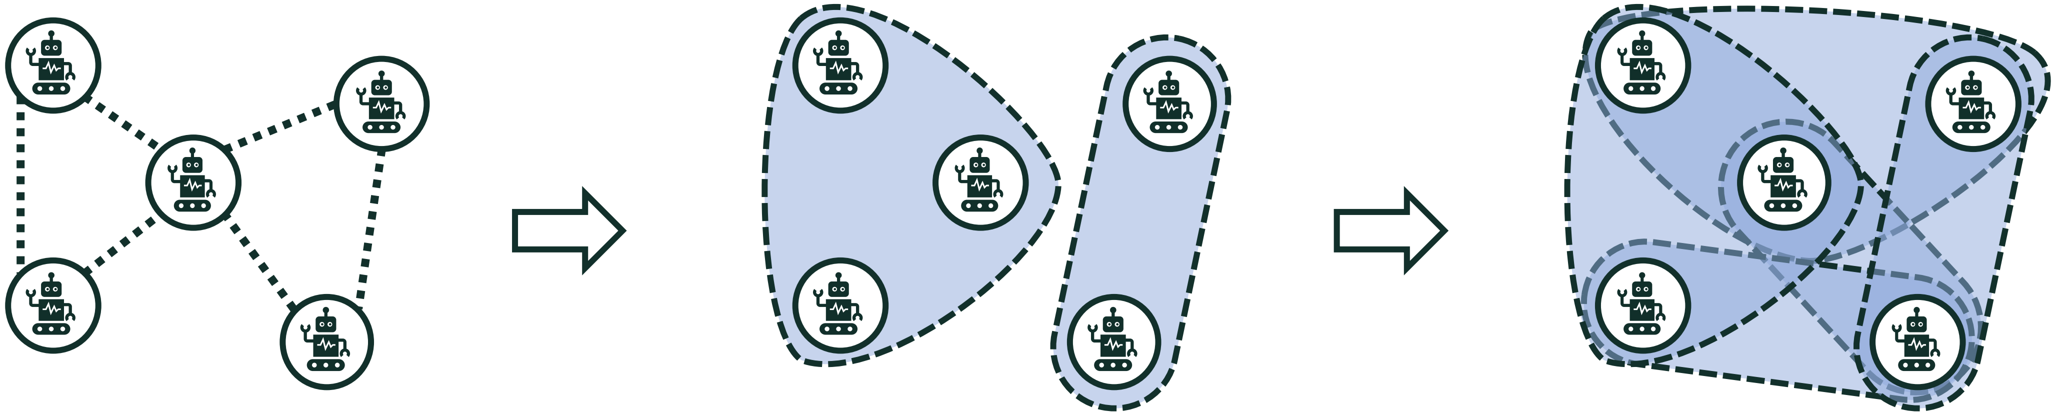

Communication Graph

- Channels over which agents can coordinate actions
- May vary with time and state of the system

Collaboration Partition

- Group of connected agents use coordinated control action
- Groups require communication and can change with time

Equilibrium Hyper-graph

- All possible groups of collaborating agents
- In equilibrium, no admissible group will deviate action

What Makes Collaboration Good?

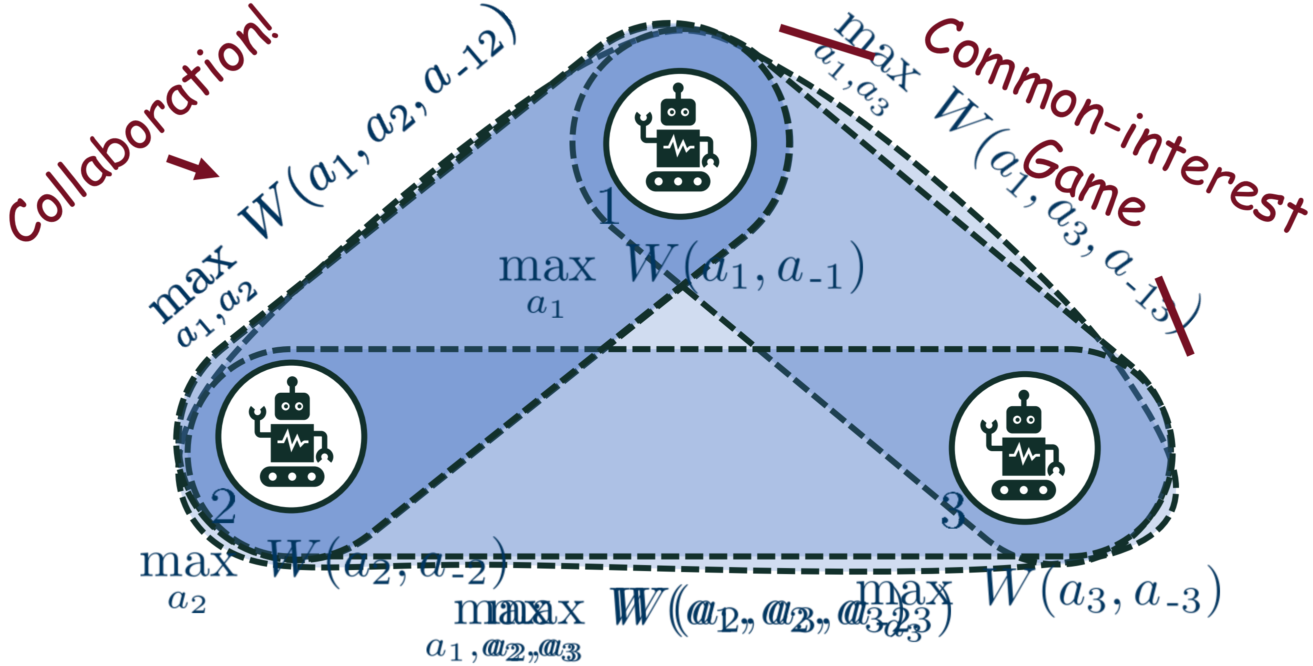

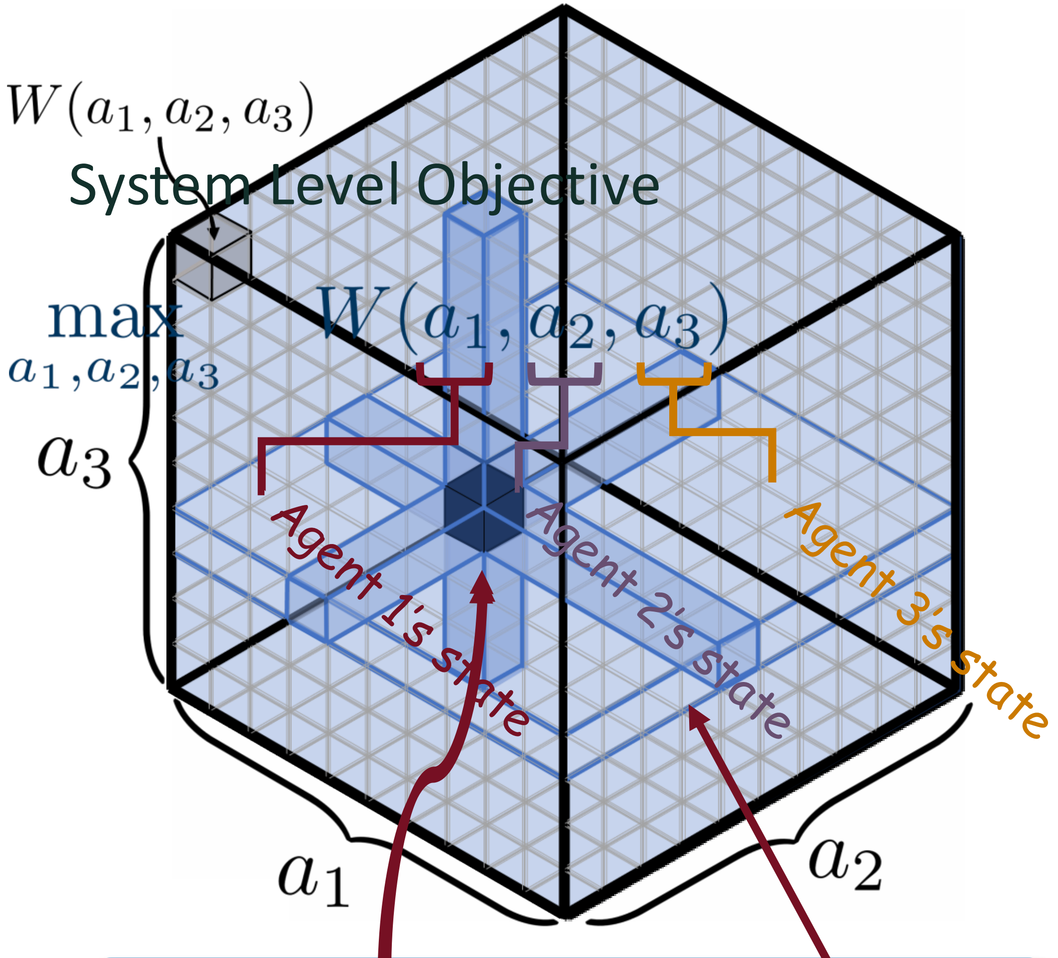

Nash Equilibrium (a^{NE}): / **2-Strong Nash Equilibrium (a^{2SNE}):**

$$W(a^{NE}) \geq W(a_i', a_{-i}^{NE})$$

$$\forall a_i' \in \mathcal{A}_i,\quad i \in N = \{1, \dots, n\}$$

$$W(a^{2SNE}) \geq W(a_\Gamma', a_{-\Gamma}^{2SNE})$$

$$\forall a_\Gamma' \in \prod \mathcal{A}_i,$$

$$\Gamma \in \{Z \subseteq N : |Z| \leq k\}$$

$\mathbf{1 \leq k \leq n}$

Bridge the gap between ***centralized*** and ***distributed*** equilibrium concepts

Collaborative Decision-Making Paradigms

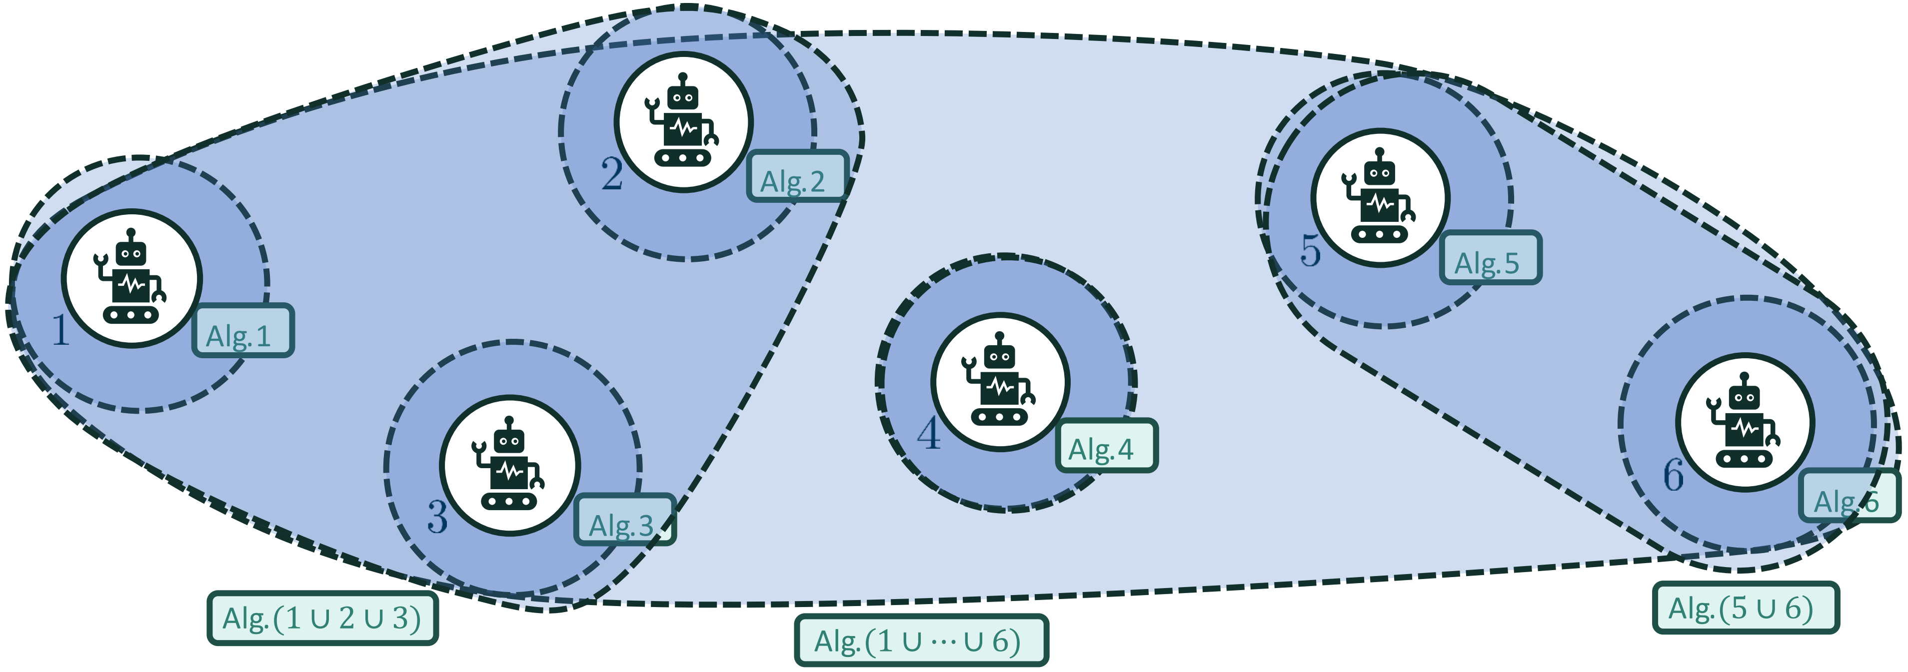

Need: system architecture to govern behavior of the agents

Distributed Decision-making
(Sub-optimal performance/ low complexity)

Partially Collaborative
Decision-making

Centralized Decision-making
(Best performance / high complexity)

Increased coordination through partial collaboration

Trade-off benefit to ***equilibrium performance*** and added ***algorithm complexity***

Efficacy of Collaborative Decision-Making

Distributed Decision-making — Centralized Decision-making

Collaboration Hyper-graph $\mathcal{H} \subseteq 2^N$

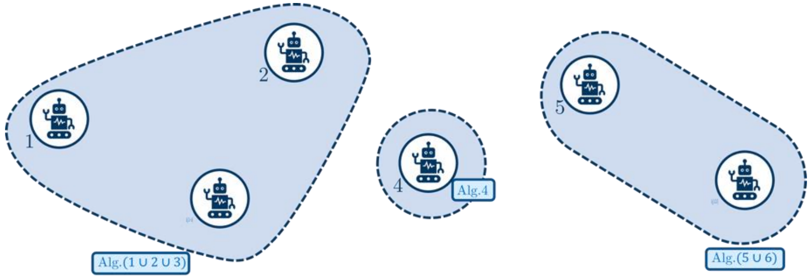

Network

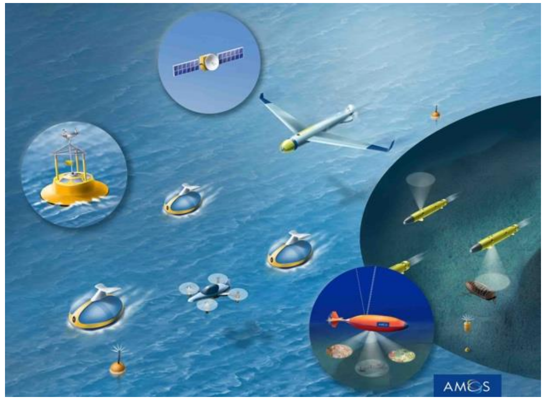

$\mathcal{H} = \{(i,j) \in N^2 \mid (i,j) \in E\}$

Intersectional

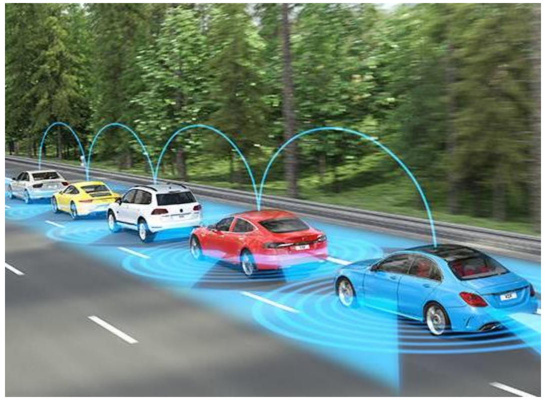

$\mathcal{H} = \{\Gamma_1, \ldots, \Gamma_z \mid \Gamma_a \cap \Gamma_{a+1} \neq \emptyset\}$

Local

$\mathcal{H} = \{\Gamma \subseteq N \mid \rho(i,j) \leq d \; \forall i,j \in \Gamma\}$

Efficacy of Collaborative Decision-Making

Distributed Decision-making — Centralized Decision-making

Collaboration Hyper-graph $\mathcal{H} \subseteq 2^N$

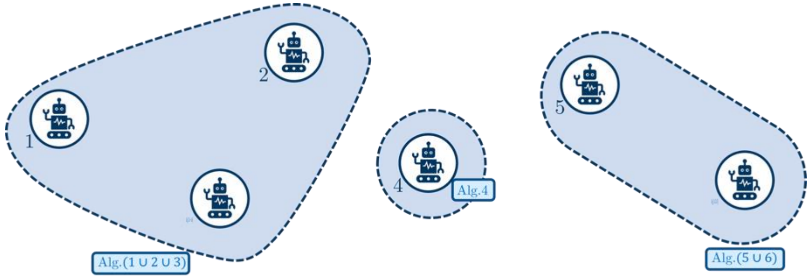

Benefit/Cost of Increased Collaboration

$$\mathcal{H}_k = \bigcup_{\zeta=1}^{k} \{\Gamma \subseteq N \mid |\Gamma| = \zeta\}$$

At most *k* agents can collaborate

k-strong Nash equilibria

As we vary the level of collaboration…

Benefit to equilibrium efficiency	Added communication costs/complexity	System Design

Coalitionally Smooth Games

How much does collaboration improve equilibrium efficiency?

Problem: Multi-Agent systems are complex and equilibria hard to compute

Definition: W is (λ, μ)-k-coalitionally-smooth, where $\lambda, \mu \in \mathbb{R}^k$, if for all $a, a' \in \mathcal{A}$

$$\frac{1}{\binom{n}{\zeta}} \sum_{\Gamma \in \mathcal{H}_\zeta} W(a'_\Gamma, a_{-\Gamma}) \geq \lambda_\zeta W(a') - \mu_\zeta W(a), \ \forall \zeta \in [k].$$

Average group deviation effect on welfare

Theorem 1.1: [Ferguson, Paccagnan, Pradelski, Marden TAC*]

If W is (λ, μ)-k-coalitionally-smooth then for each k-strong Nash equilibrium $a^{k\mathrm{SNE}}$

$$\mathbf{Eff}(W) := \frac{W(a^{k\mathrm{SNE}})}{\max_{a \in \mathcal{A}} W(a)} \geq \frac{\lambda_\zeta}{1 + \mu_\zeta}, \ \forall \zeta \in [k].$$

For any W there exists some λ and μ that satisfy the definition

Equilibrium Welfare / Optimal Welfare

Measure of equilibrium efficiency

0 - very inefficient

1 - fully efficient

Just need to find the right ones...

Complicated Multi-Agent Systems Analysis → Two Parameter Search (λ, μ)

Resource Allocation Problems

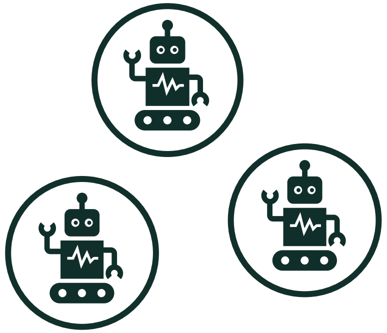

Agents (N)

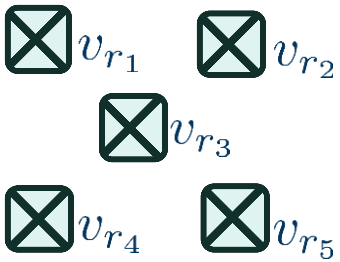

Resources $(\mathcal{R})$

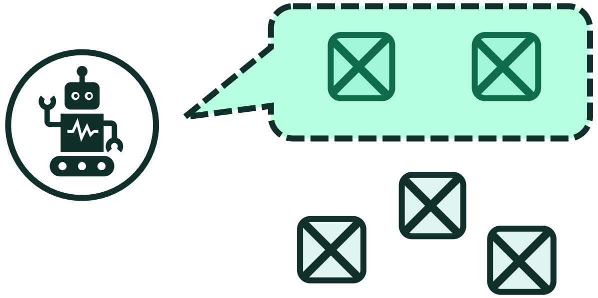

Actions $(a_i \in \mathcal{A}_i \subseteq 2^{\mathcal{R}})$

Resource Allocation Problems

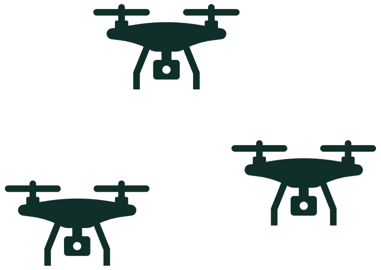

Agents (N)

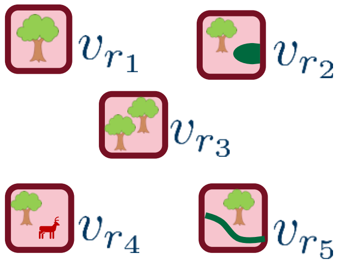

Resources $(\mathcal{R})$

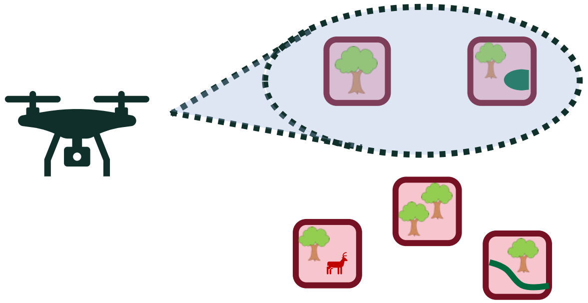

Actions $(a_i \in \mathcal{A}_i \subseteq 2^{\mathcal{R}})$

Drone Surveillance

Resource Allocation Problems

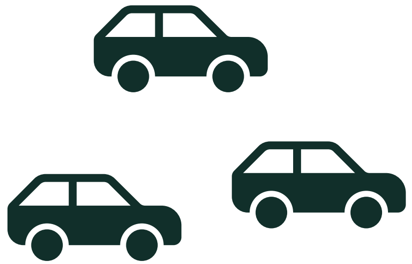

Agents (N)

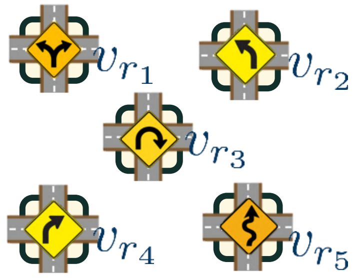

Resources $(\mathcal{R})$

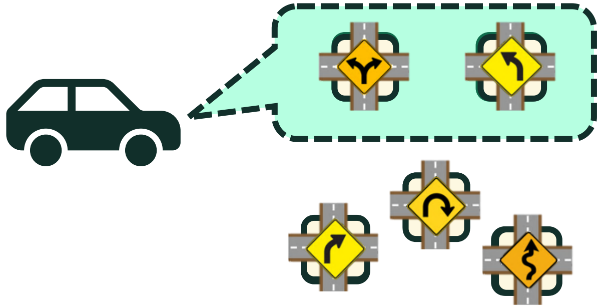

Actions $(a_i \in \mathcal{A}_i \subseteq 2^{\mathcal{R}})$

Traffic Routing

Resource Allocation Problems

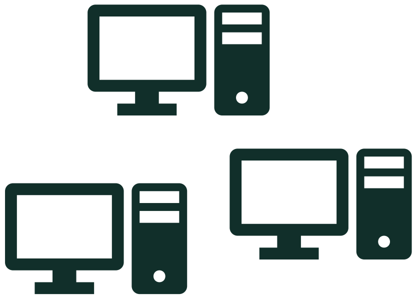

Agents (N)

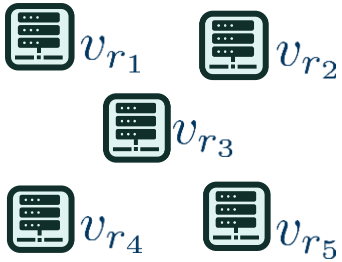

Resources $(\mathcal{R})$

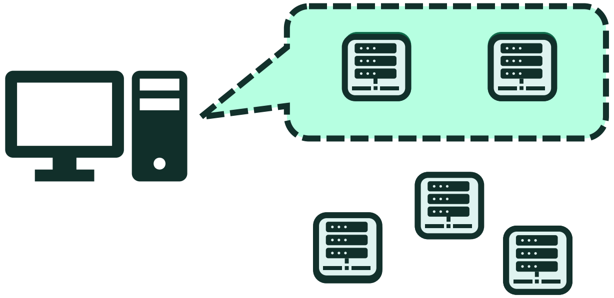

Actions $(a_i \in \mathcal{A}_i \subseteq 2^{\mathcal{R}})$

Cloud Computing/Internet Routing

Resource Allocation Problems

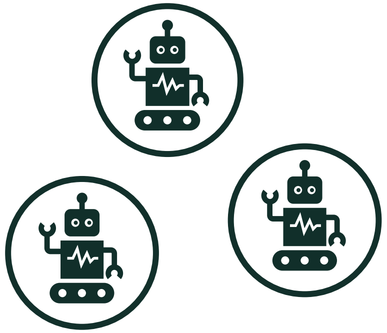

Agents (N)

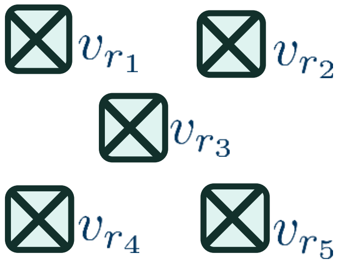

Resources $(\mathcal{R})$

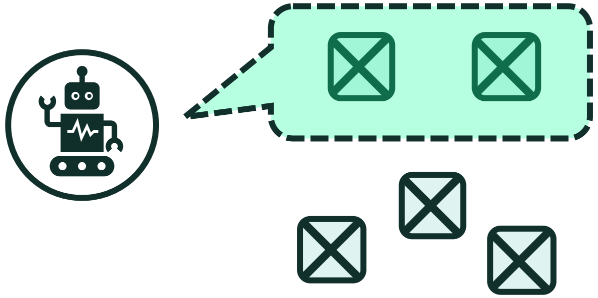

Actions $(a_i \subseteq \mathcal{R})$

System Welfare

$$W(a) = \sum_{r \in \mathcal{R}} v_r \underbrace{w(|a|_r)}$$

Value of having $|a|_r$ agents share resource r

Diminishing returns

Number of agents on resource: $|a|_r$

Centralized optimal solution: NP-hard

Partially collaborative equilibria: $\texttt{Eff}(k)$

Distributed equilibria: known sub-optimal performance

Quantifying k-Strong Efficiency

Theorem 1.1: [Ferguson, Paccagnan, Pradelski, Marden TAC*]

Find (λ, μ) for bound on $\texttt{Eff}(k)$

Proposition 1.2: [Ferguson, Paccagnan, Pradelski, Marden TAC*]

Any resource allocation game with welfare function w is $\lambda_\zeta = 1/\nu^\star_\zeta$ and $\mu_\zeta = \rho^\star_\zeta/\nu^\star_\zeta - 1$, where $(\rho^\star_\zeta, \nu^\star_\zeta)$ is a solution to the linear program:

(λ, μ) are the solution to a tractable linear program

$$(\rho^\star_\zeta, \nu^\star_\zeta) \in \arg\min_{\rho,\nu \geq 0} \quad \rho$$

$$\text{s.t.} \qquad 0 \geq w(o+x) - \rho$$

Example: Covering Problems

$w(|a|_r) = \mathbb{1}\,[|a|_r \geq 1]$

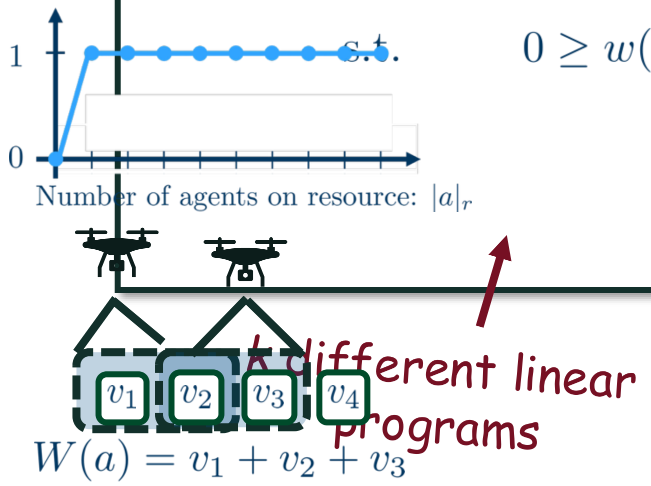

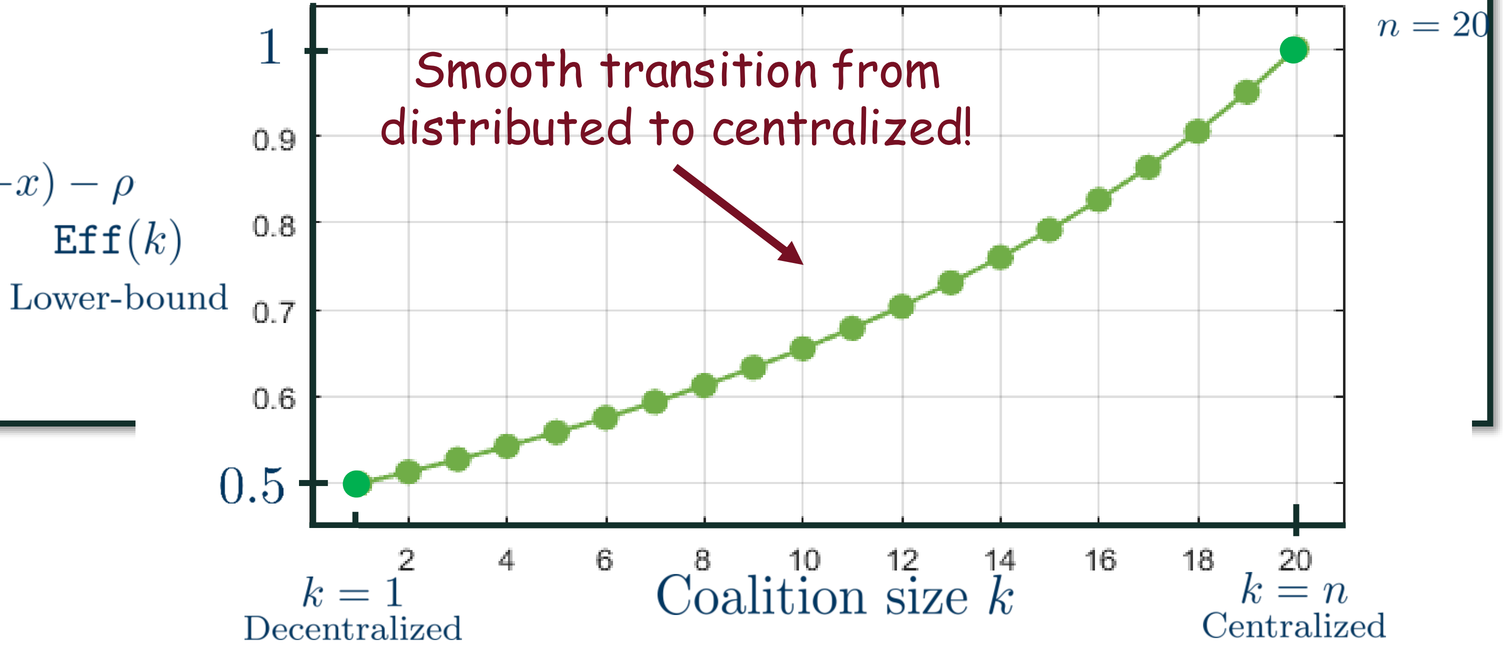

Efficacy of Collaborative Decision-Making

Distributed Decision-making

Centralized Decision-making

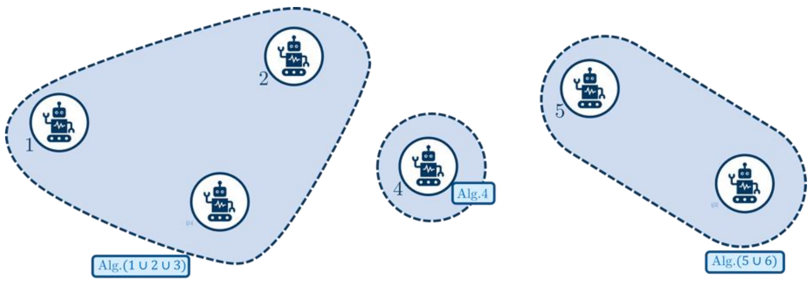

Benefit/Cost of Increased Collaboration

As we vary the level of collaboration…

Benefit to equilibrium efficiency

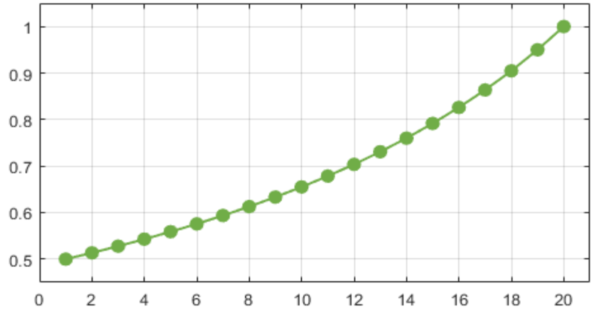

Added communication costs/complexity

System Design

Complexity of *k*SNE

How does collaboration affect ***Computational Complexity***?

$n :=$ number of players $\quad m :=$ number of agent actions $\quad k :=$ size of coalitions

	Nash Equilibrium	k-Strong Nash Equilibrium
Evaluating an equilibrium	$\mathcal{O}(nm)$	$\mathcal{O}\left(\frac{n!}{(n-k)!k!}m^k\right)$

k-Round-Robin: each group updates in set order

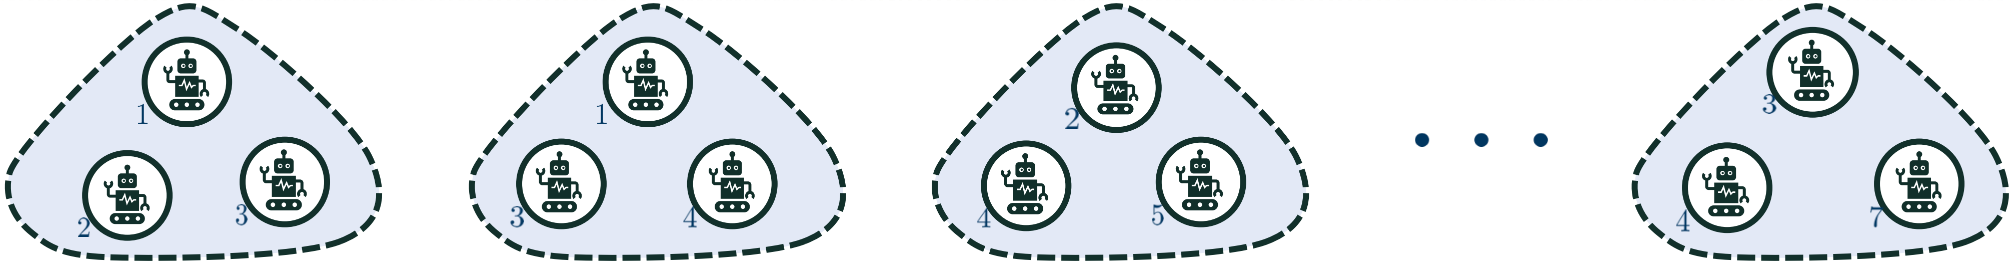

Proposition 1.3: [Ferguson, Paccagnan, Pradelski, Marden TAC*]

The k-Round-Robin dynamics converge in finite time and requires welfare comparisons on order $\mathcal{O}\left(m^n\left(\frac{1}{1-\frac{1}{m}}\right)^k\right)$

Asynchronous Dynamics

What is the performance like ***before*** convergence?

k-Asynchronous-Best-Response: each group updates in a ***random*** order

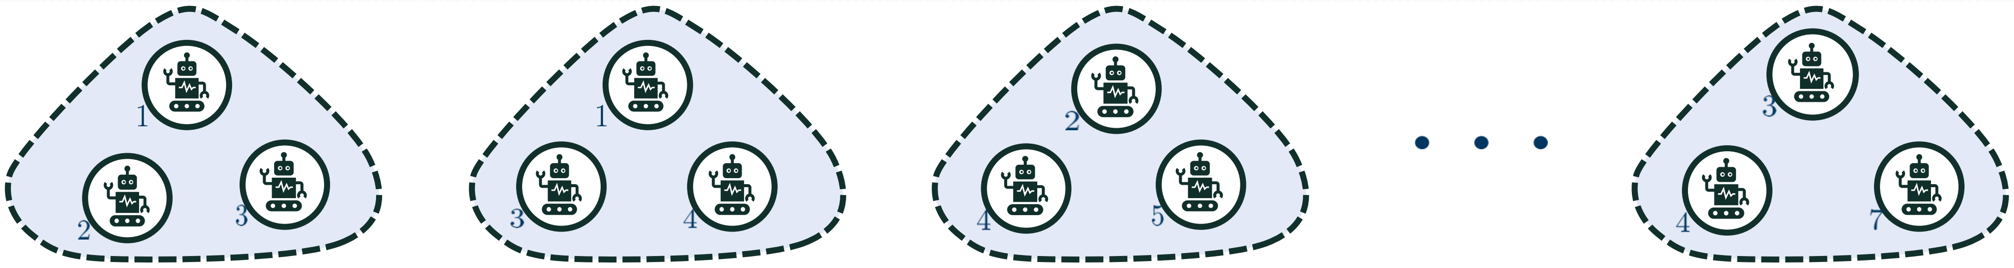

Theorem 1.4: [Ferguson, Paccagnan, Pradelski, Marden TAC*]

If W is (λ, μ)-k-coalitionally smooth, then after $T \geq 1$ update steps, the cumulative expected welfare satisfies

$$\mathbb{E}\left[\frac{1}{T}\sum_{t=1}^{T} W(a^t)\right] \geq \frac{T-1}{2T}\,\frac{\sum_{\zeta=1}^{k} p_\zeta \lambda_\zeta}{1+\sum_{\zeta=1}^{k} p_\zeta \mu_\zeta}\,W(a^{\text{opt}}),$$

where p_ζ is the probability a group of size ζ best responds.

Sample bigger groups
Better guarantee!

Efficacy of Collaborative Decision-Making

Distributed
Decision-making

Centralized
Decision-making

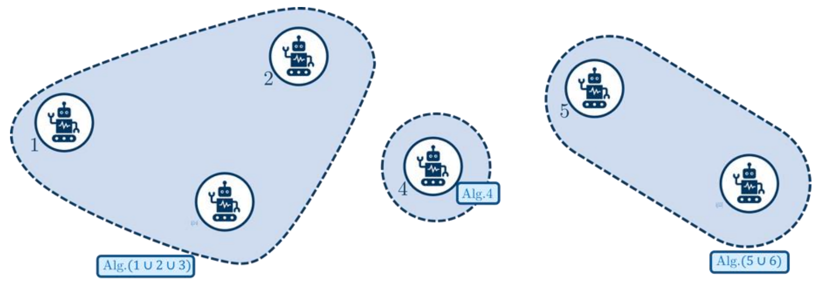

Benefit/Cost of Increased Collaboration

As we vary the level of collaboration…

Benefit to equilibrium efficiency

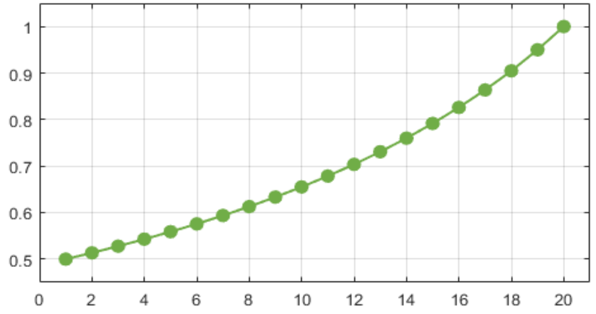

Added communication costs/complexity

System Design

Utility Design for Collaborative Systems

System Objective

$$\max_{a_1, a_2, \ldots, a_n} W(a)$$

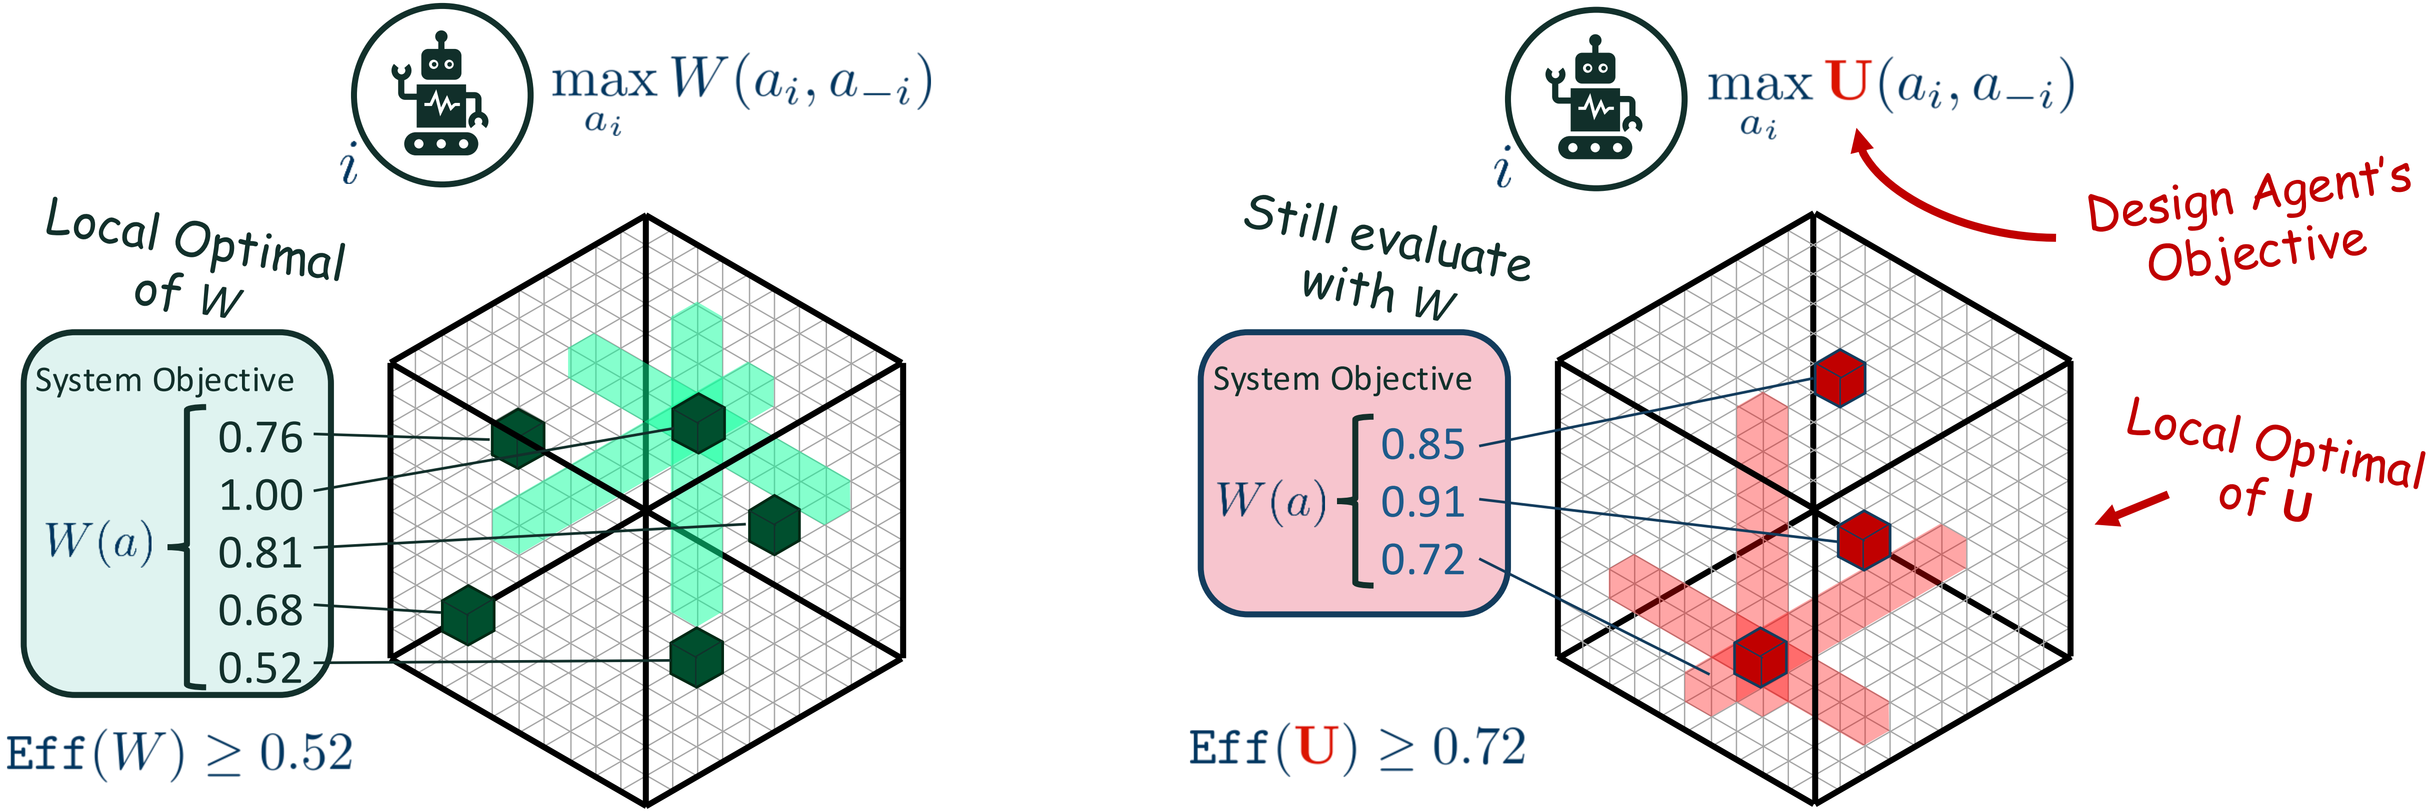

Set of ... N... quilibria ... $(\mathbf{U})$

How do we design the agent objective function?

Can we do this for the collaborative setting?

Utility Design for Collaborative Systems

Design to optimize worst-case equilibrium performance

Utility Design Problem

$$\max_{\mathbf{U}:\mathcal{A}\to\mathbb{R}} \frac{\min_{a^{k\mathrm{SNE}}\in k\mathrm{SNE}(\mathbf{U})} W(a^{k\mathrm{SNE}})}{\max_{a^{\mathrm{opt}}\in\mathcal{A}} W(a^{\mathrm{opt}})}$$
$$\text{s.t.} \quad \text{informational constraints}$$

with respect to original objective!

Theorem 1.5: [Ferguson, Paccagnan, Pradelski, Marden TAC*]

In resource allocation problems, we can solve the utility design problem for *k*-strong Nash equilibria by extending the prior linear programming technique

Utility Design for Collaborative Systems

Design to optimize worst-case equilibrium performance

Utility Design Problem

$$\max_{\mathbf{U}:\mathcal{A}\to\mathbb{R}} \frac{\min_{a^{k\mathrm{SNE}}\in k\mathrm{SNE}(\mathbf{U})} W(a^{k\mathrm{SNE}})}{\max_{a^{\mathrm{opt}}\in\mathcal{A}} W(a^{\mathrm{opt}})}$$

$$\text{s.t.} \quad \text{informational constraints}$$

with respect to original objective!

In Resource Allocation Problems,

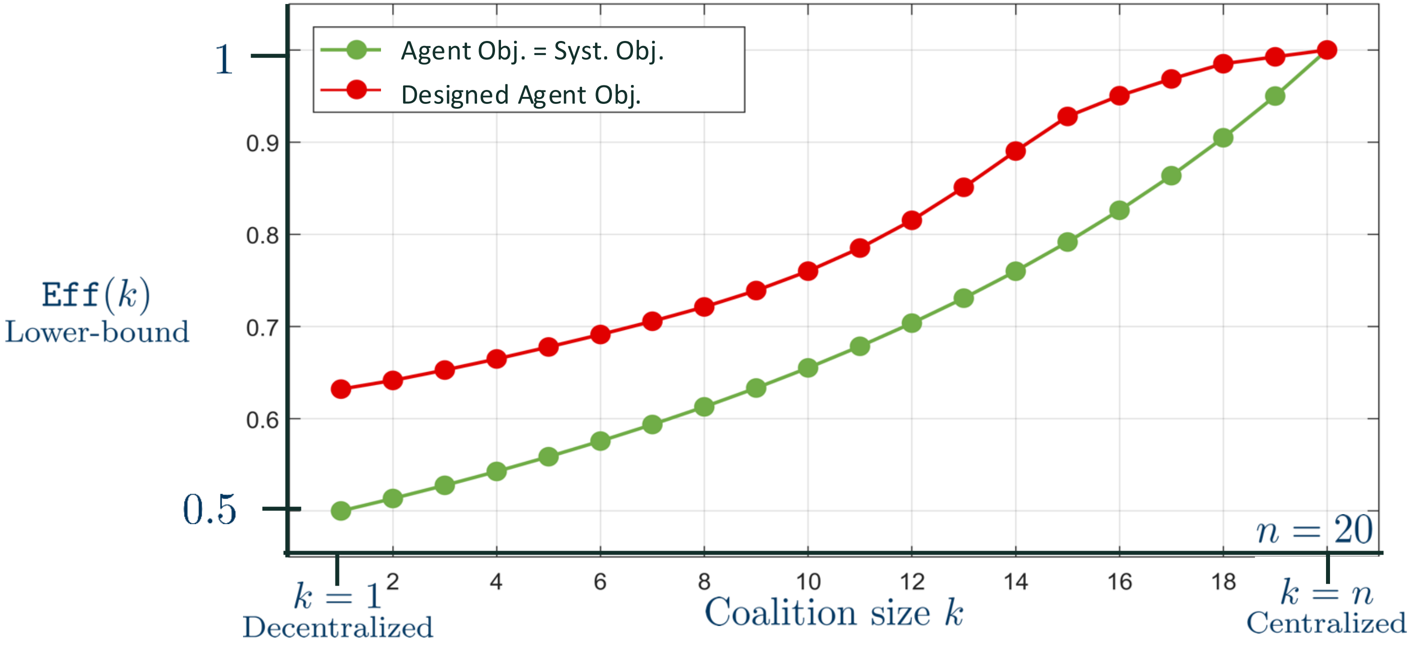

In **designing collaborative systems**, consider:

1) The ***level of collaboration***

2) How ***collaborative decisions*** are made

We find significant opportunities for improvement

Efficacy of Collaborative Decision-Making

Distributed
Decision-making

Centralized
Decision-making

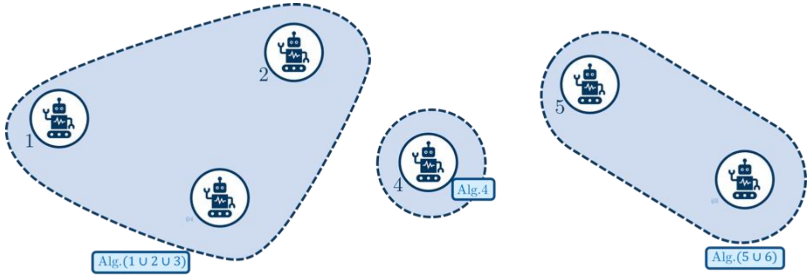

Benefit/Cost of Increased Collaboration

As we vary the level of collaboration…

Benefit to equilibrium efficiency

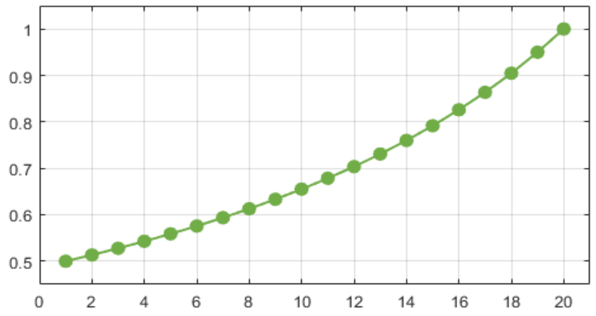

Added communication costs/complexity

System Design

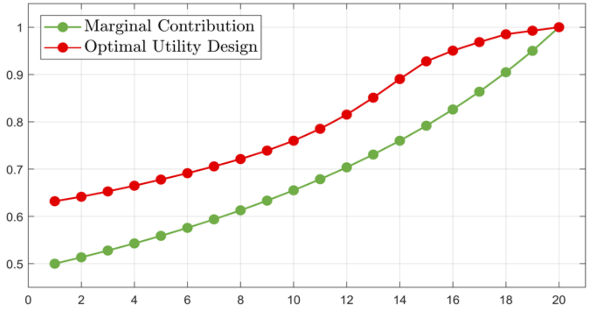

Engineered Multi-Agent Systems

Drone Fleets

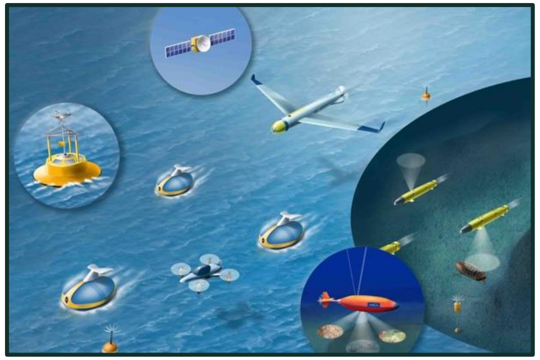

Manufacturing Robotics

Internet of Things

Large, interconnected systems of many automated devices and processes

Quantifying the effects of Collaborative Decision-Making

[CDC23, TAC*, CDC24]

1) Extend game theoretic approach to collaborative settings
2) Gain understanding of the benefits and costs of collaboration
3) New insights on how to design collaborative system

Distributed Control
(local decisions)

Design new system paradigms
between centralized and distributed

Centralized Control
(Coordinated decisions)

Further Directions

Broader scope of inter-agent communication and coordination

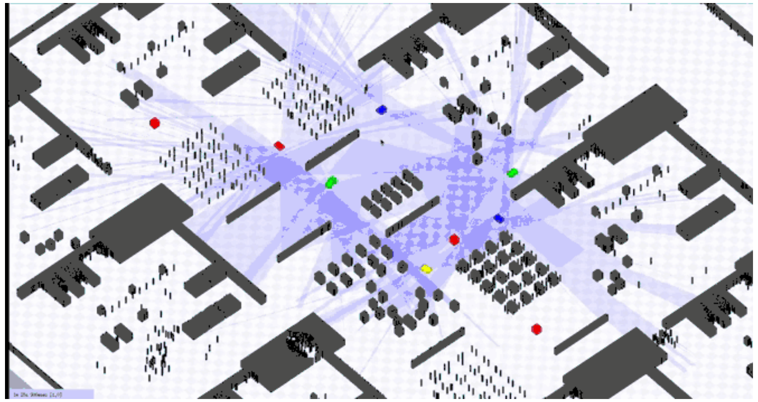

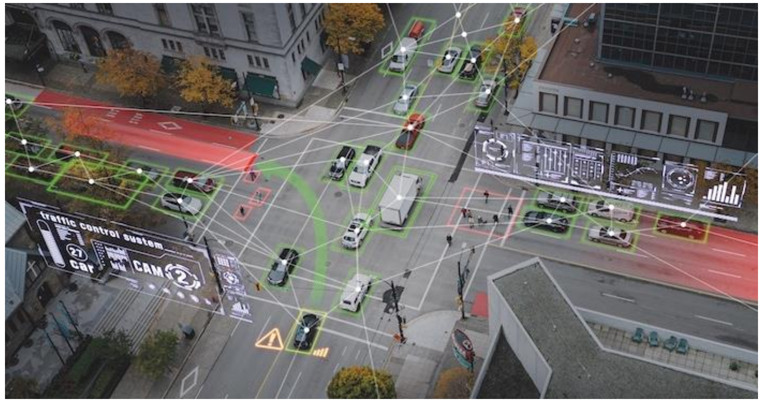

Local and Dynamic Connectivity

- Evolving Objectives
- Online Task Assignment
- Dynamic Team Formation

Multi-Agent Learning

- Collection of Learners
- Explore vs Exploit vs Comm.
- Collaborative Learning Process

Data and Distributed Estimation

- Relevant Info. Beyond Collab.
- Estimate using shared data
- Selective Information Sharing

II. Inter-agent Competition

Designing Competitive Controllers

Design Local Control Law

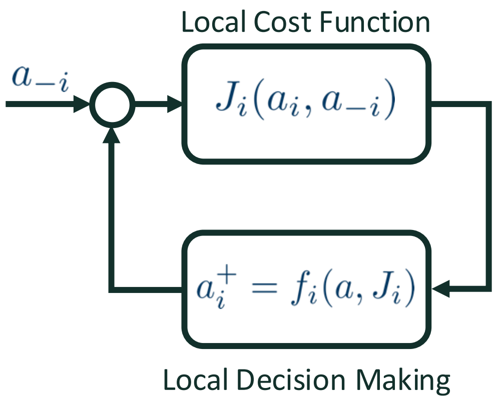

Converge to Nash Equilibria

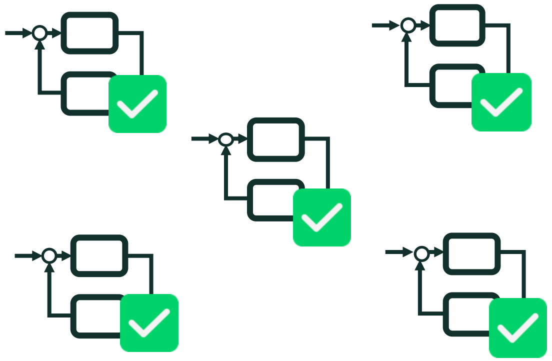

Quality of Nash Equilibria

$$\texttt{Eff} = \frac{W(a^{\mathrm{NE}})}{W(a^{\mathrm{opt}})}$$

Designing Competitive Controllers

Design ***Competitive*** Control Law

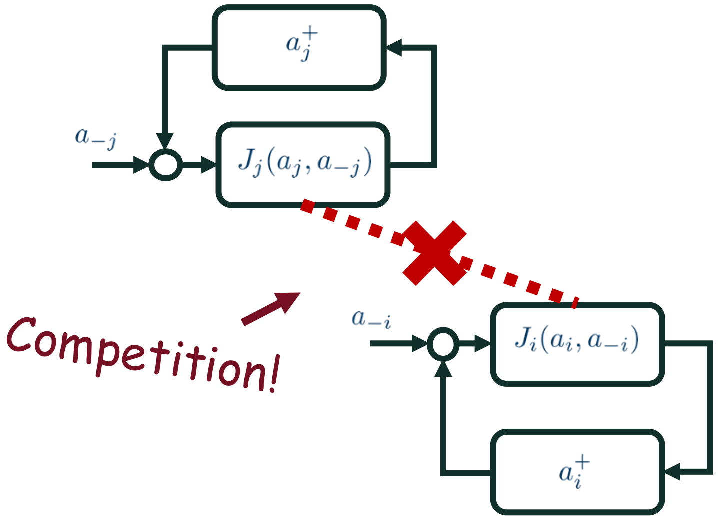

Competitive control design

Converge to Nash Equilibria

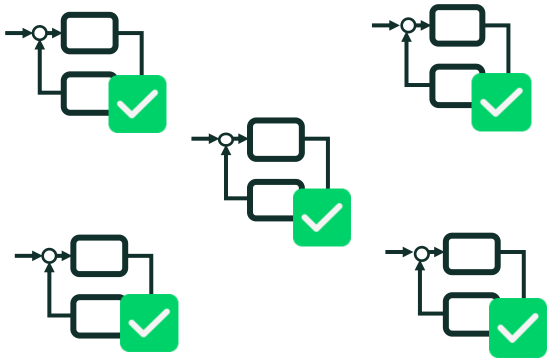

Quality of Nash Equilibria

$$\texttt{Eff} = \frac{W(a^{\mathrm{NE}})}{W(a^{\mathrm{opt}})}$$

Designing Competitive Controllers

Design ***Competitive*** Control Law

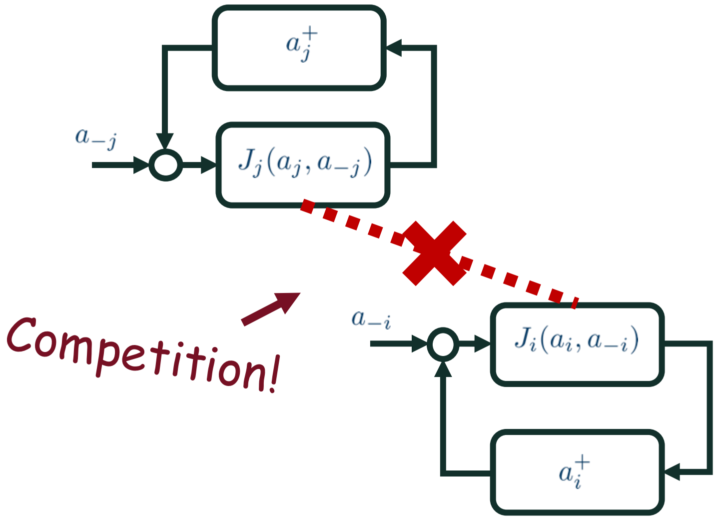

Competitive control design

Predict and steer emergent behavior

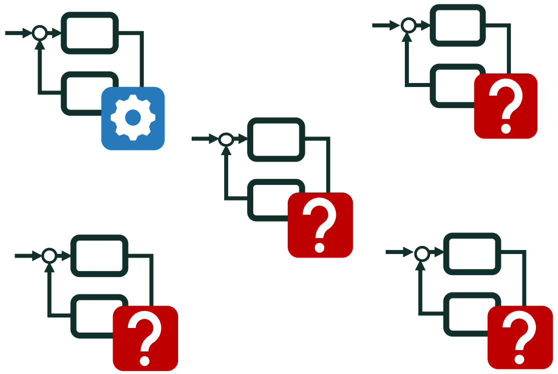

Accuracy of Predicted Outcome

$$J_i(a^{\text{predicted}}) \text{ vs } J_i(a^{\text{realized}})$$

Ego perspective on prediction accuracy

Designing Competitive Controllers

Design ***Competitive*** Control Law

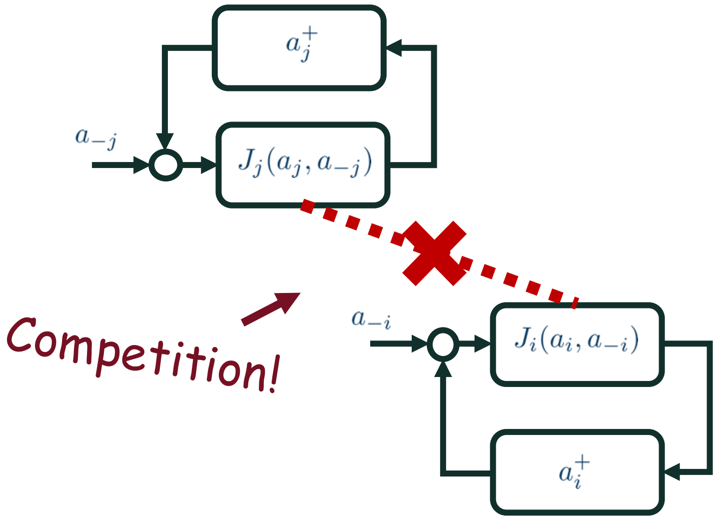

Competitive control design

Predict and steer emergent behavior

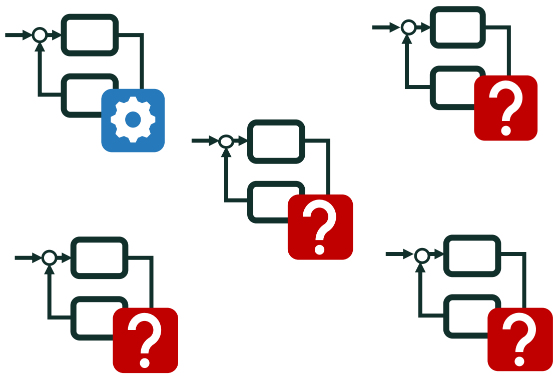

Accuracy of Predicted Outcome

$$J_i(a^{\text{predicted}}) \text{ vs } J_i(a^{\text{realized}})$$

Ego perspective on prediction accuracy

When agents compete, how does the ***accuracy of their prediction*** of other agents' behavior affect the collective outcome

A Motivating Example: Autonomous Racing

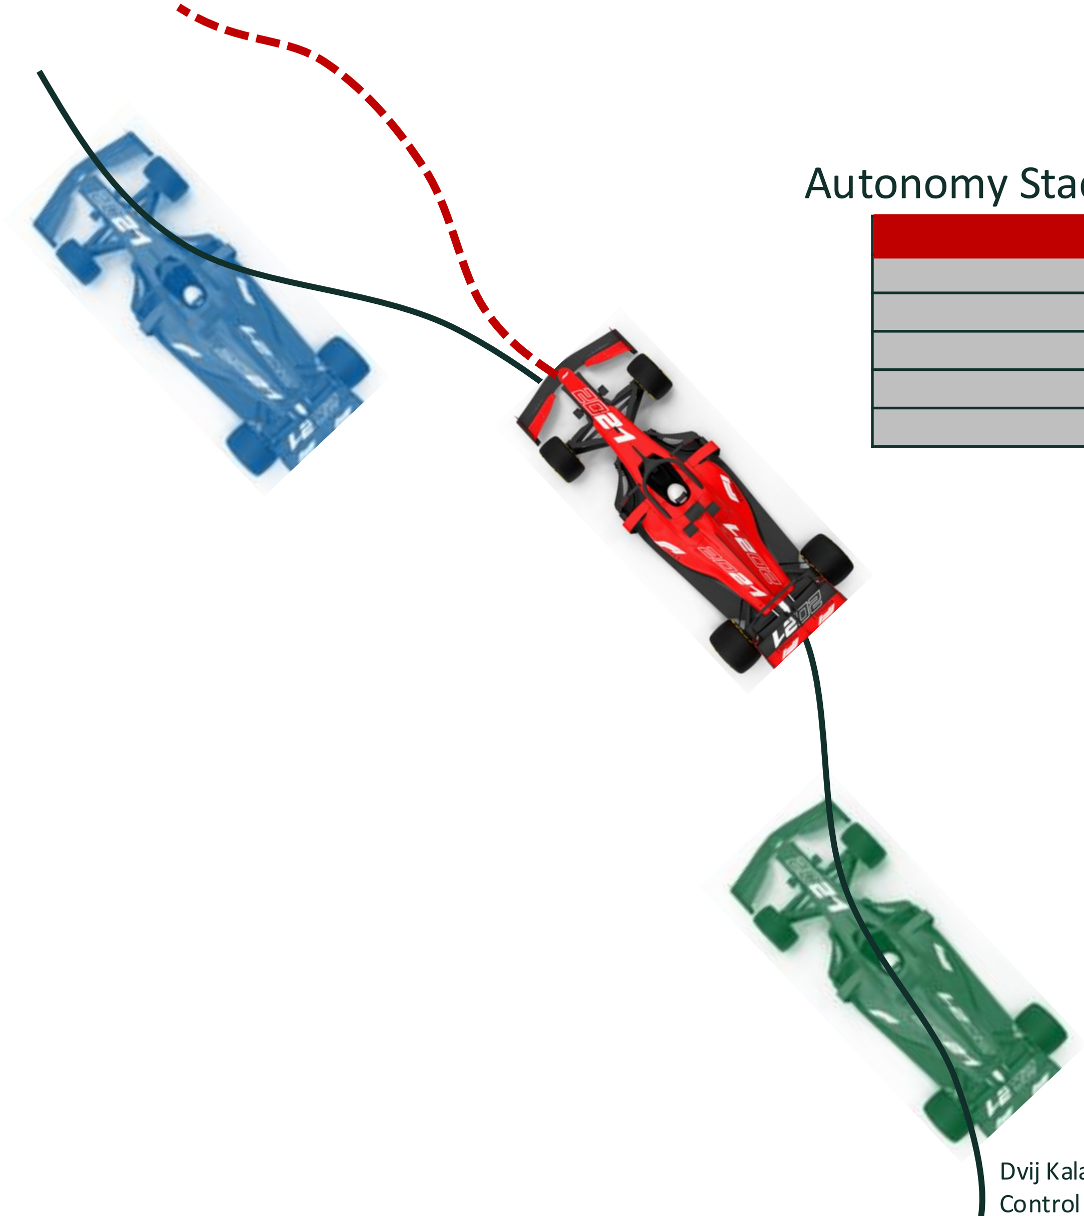

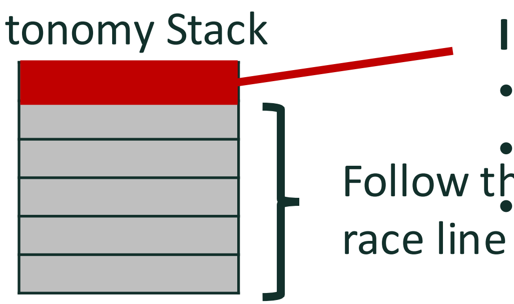

Intervehicle strategy

- Passing
- Blocking
- Using limited fuel

Approach:

1. Design rewards
   $$r_i(u, \theta) = \text{dist ahead or behind}$$
2. Cast as a large Markov Game
   $$J_i(\pi, s_0) = \mathbb{E}_{u^t \sim \pi(s^t)} \left[\sum_{t=0}^{\infty} \gamma^t r_i(u^t, s^t)\right]$$
3. Solve for a Nash equilibrium joint policy
   π^{NE} via Actor-Critic MARL
4. Take ego policy and implement in the autonomy stack
   $$\pi^{\text{NE}} = (\pi_1^{\text{NE}}, \ldots, \pi_{ego}^{\text{NE}}, \ldots, \pi_n^{\text{NE}})$$

Dvij Kalaria, Chinmay Maheshwari, and Shankar Sastry. "{\alpha}-RACER: Real-Time Algorithm for Game-Theoretic Motion Planning and Control in Autonomous Racing using Near-Potential Function." *arXiv preprint arXiv:2412.08855*(2024).

A Motivating Example: Autonomous Racing

Simulated Multi-Agent Interaction

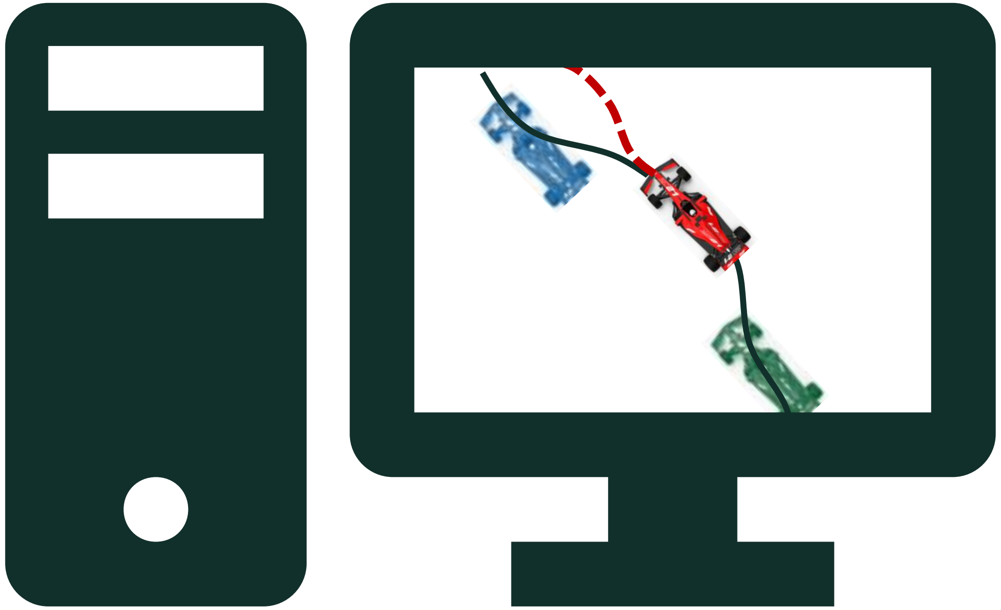

Synthesized Policy

π_{ego}^{NE}

Realized Multi-Agent Interaction

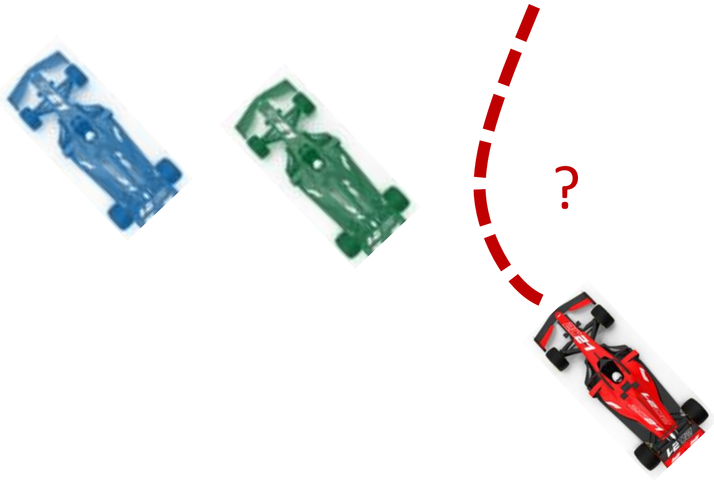

Difference between ***simulated*** and ***realized*** interaction can cause degraded performance of policy

How do we quantify and overcome this ***'Strategic' sim-to-real gap?***

Realizations of Game Theoretic Planning and Prediction

Dynamic Zero-Sum Games

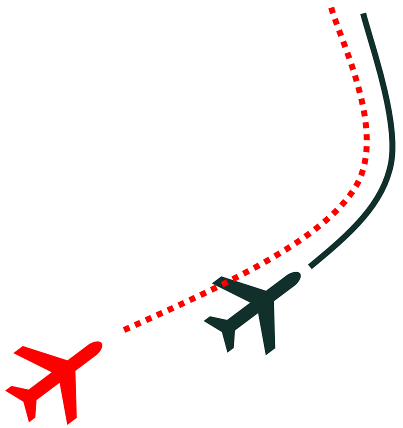

Iterative LQ

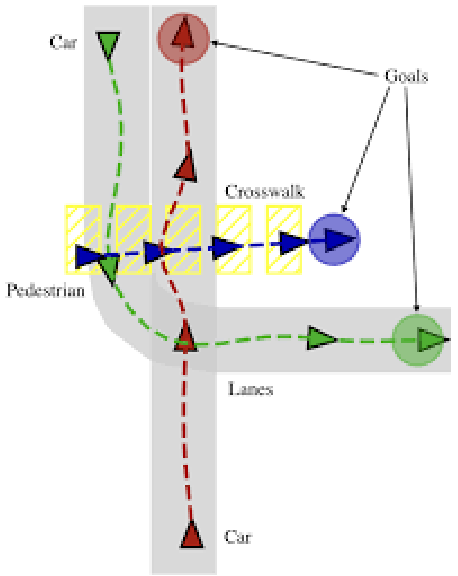

Model Predictive Games

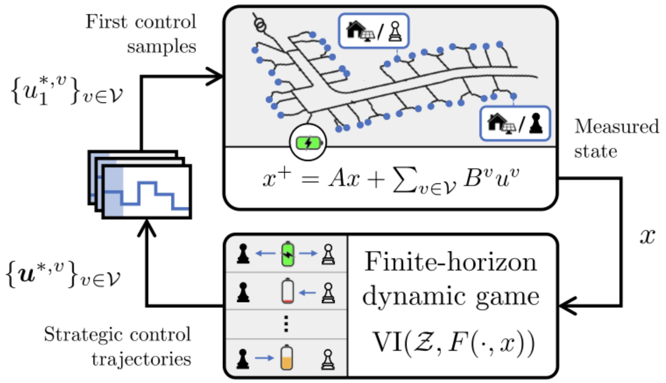

1. J. Sprinkle, J. M. Eklund, H. J. Kim and S. Sastry, "Encoding aerial pursuit/evasion games with fixed wing aircraft into a nonlinear model predictive tracking controller," 2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601), Nassau, Bahamas, 2004, pp. 2609-2614 Vol.3, doi: 10.1109/CDC.2004.1428851.
2. Başar, Tamer, ed. *Dynamic games and applications in economics*. Vol. 265. Springer Science & Business Media, 1986.
3. H. Chen, C. W. Scherer and F. Allgower, "A game theoretic approach to nonlinear robust receding horizon control of constrained systems," *Proceedings of the 1997 American Control Conference (Cat. No.97CH36041)*, Albuquerque, NM, USA, 1997, pp. 3073-3077 vol.5, doi: 10.1109/ACC.1997.612023.
4. D. Fridovich-Keil, E. Ratner, L. Peters, A. D. Dragan and C. J. Tomlin, "Efficient Iterative Linear-Quadratic Approximations for Nonlinear Multi-Player General-Sum Differential Games," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, Paris, France, 2020
5. S. Hall, G. Belgioioso, D. Liao-McPherson and F. Dorfler, "Receding Horizon Games with Coupling Constraints for Demand-Side Management," 2022 IEEE 61st Conference on Decision and Control (CDC), Cancun, Mexico, 2022, pp. 3795-3800, doi: 10.1109/CDC51059.2022.9992497.
6. E. Benenati and G. Belgioioso. "The explicit game-theoretic linear quadratic regulator for constrained multi-agent systems." arXiv preprint arXiv:2512.07749 (2025).

Games as Models

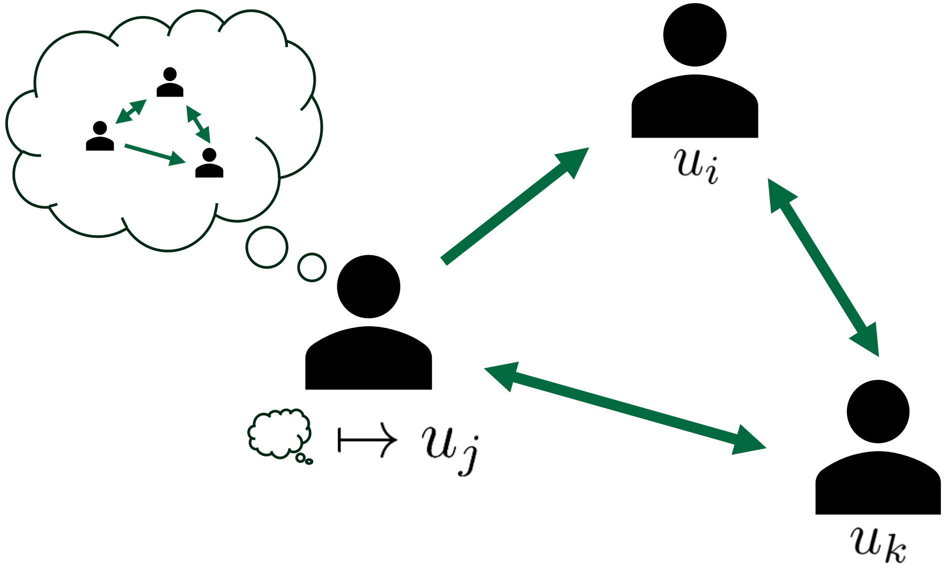

- Each agent possesses a model for how the collective group will interact
- Use conjectures about other players' intents and capabilities to predict the outcome and your role in it
- Opt for models that possess efficiently computed solutions

Games as Models

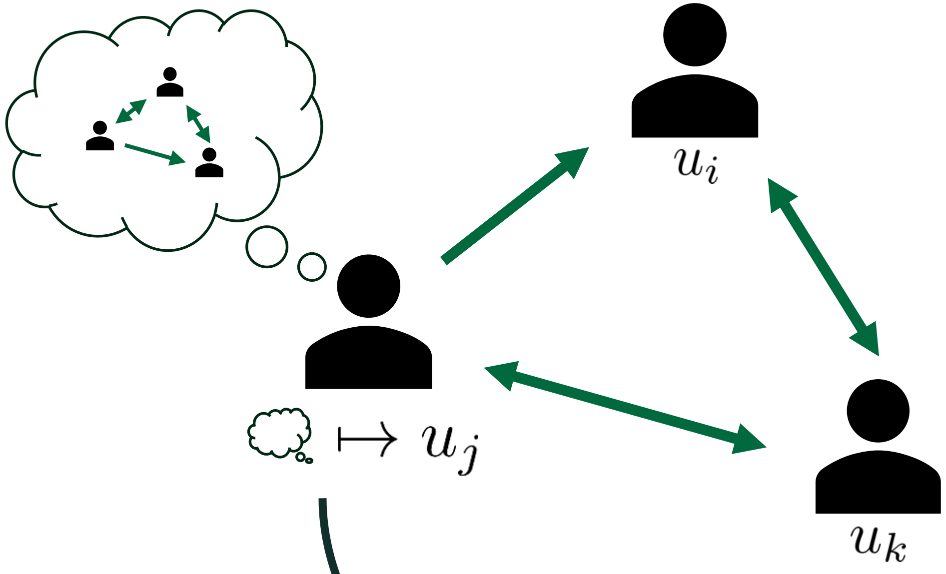

- Each agent possesses a model for how the collective group will interact
- Use conjectures about other players' intents and capabilities to predict the outcome and your role in it
- Opt for models that possess efficiently computed solutions

Games and Prediction

$$G = \langle N, \{\mathcal{U}_i\}_{i \in N}, \{J_i\}_{i \in N}, \texttt{Soln} \rangle$$

Players — Control/Action set — Objective — Solution Concept

$$\texttt{Soln} : \{J : \mathcal{U} \to \mathbb{R}\}^n \to 2^{\mathcal{U}} \qquad \texttt{Soln}\left(\{J_i\}_{i \in N}\right) \subseteq \mathcal{U}$$

Predictive Capabilities of Games

$$G^{(j)} = \langle N, \{\mathcal{U}_i^{(j)}\}_{i\in N}, \{J_i^{(j)}\}_{i\in N}, \texttt{Soln}^{(j)}\rangle$$

Conjectured game of player *j*

$$\text{Player } j \text{ predicted behavior: } u^{(j)} \in \texttt{Soln}^{(j)}(J^{(j)})$$

Predictive Capabilities of Games

$$G^{(j)} = \langle N, \{\mathcal{U}_i^{(j)}\}_{i \in N}, \{J_i^{(j)}\}_{i \in N}, \texttt{Soln}^{(j)} \rangle$$

Conjectured game of player j

$$\text{Player } j \text{ predicted behavior: } u^{(j)} \in \texttt{Soln}^{(j)}(J^{(j)})$$

Each player predicts behavior individually

$G^{(1)}$ $G^{(2)}$ $\cdots$ $G^{(j)}$ $\cdots$ $G^{(n)}$

Players act in response to their own predictions

Realized Behavior: $u^{\circ} = (u_1^{(1)}, u_2^{(2)}, \ldots, u_j^{(j)}, \ldots, u_n^{(n)})$

- Need not be an equilibrium/solution
- Gap between simulated and real behavior
- Exaggerated by greater mischaracterizations

Misspecified Game-theoretic Planning

Players conjecture different games

$$G^{(i)} \neq G^{(j)}$$

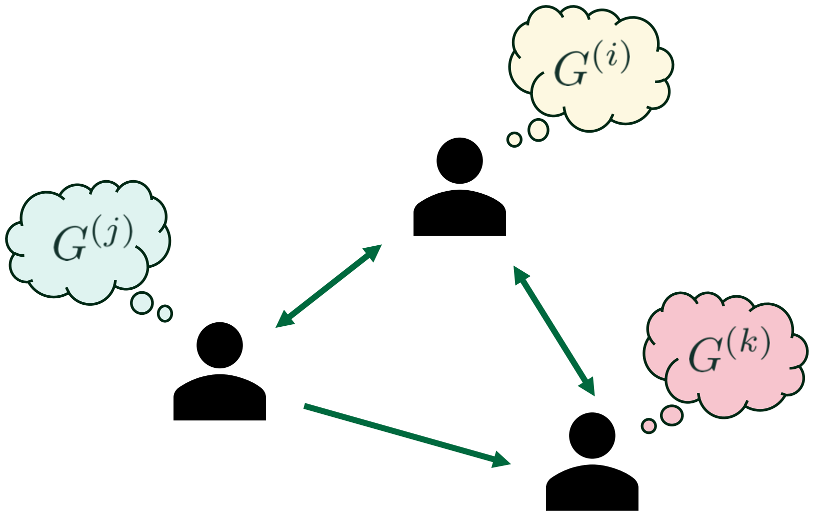

Misconjectured cost functions $J_i \neq J_i^{(j)}$

Misconjectured constraints $\mathcal{U}_i \neq \mathcal{U}_i^{(j)}$

Misaligned solution concepts $\texttt{Soln}^{(i)} \neq \texttt{Soln}^{(j)}$

Multiplicity of equilibria/predictions $|\texttt{Soln}(G)| > 1$

Misspecified Game-theoretic Planning

Players conjecture different games

$$G^{(i)} \neq G^{(j)}$$

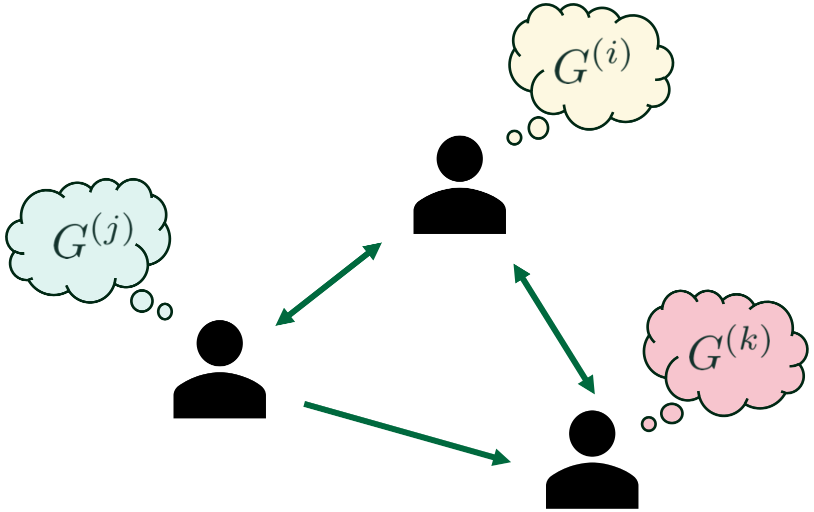

Misconjectured cost functions $J_i \neq J_i^{(j)}$

Misconjectured constraints $\mathcal{U}_i \neq \mathcal{U}_i^{(j)}$

Misaligned solution concepts $\texttt{Soln}^{(i)} \neq \texttt{Soln}^{(j)}$

Multiplicity of equilibria/predictions $|\texttt{Soln}(G)| > 1$

Game-to-Real Gap

$$J_i\big(u^{\circ}\big) - J_i\big(u^{(i)}\big)$$

Quantify magnitude

Dynamics with misspecification

Sampling, learning, and overcoming

Related Perspectives

- Sim-to-real gap in RL
- Sensitivity of equilibria in games
- Bayesian Games
- ‘Robust’ game theory
- Strategically robust equilibria
- K-level rationality
- Inverse learning in games

Our Perspective:

Adopt classic, computationally friendly games and equilibria as interaction models, even though they may not be entirely accurate (as in practice, none will)

How much do we lose from our model being incorrect, and how do we overcome?

Recent work: *Game-to-Real Gap: Quantifying the Effect of Model Misspecification in Network Games*

- Focus on well-studied case of Nash equilibria in strongly monotone network games
- Characterize and bound the game-to-real gap conditioned on game-defining parameters
- Identify forecast and network specific properties

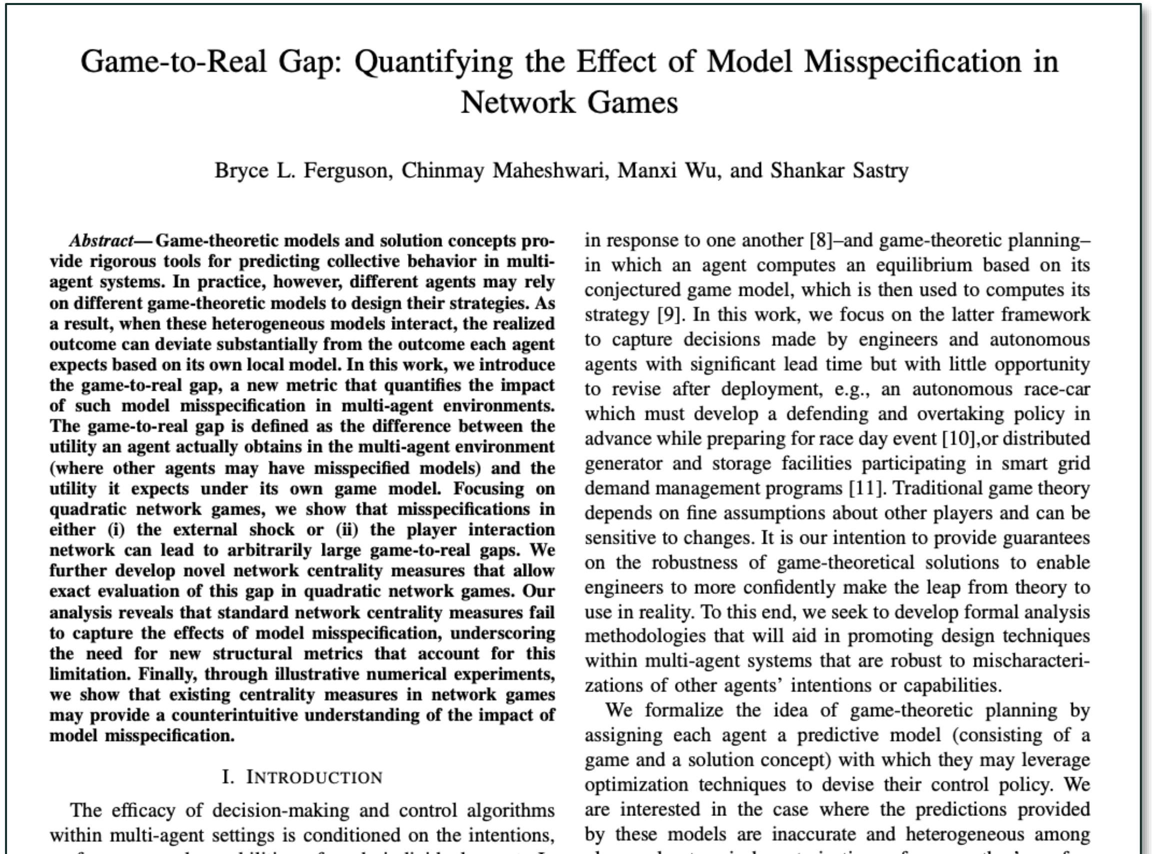

Game-to-Real Gap: Quantifying the Effect of Model Misspecification in Network Games

Bryce L. Ferguson, Chinmay Maheshwari, Manxi Wu, and Shankar Sastry

***Abstract*— Game-theoretic models and solution concepts provide rigorous tools for predicting collective behavior in multi-agent systems. In practice, however, different agents may rely on different game-theoretic models to design their strategies. As a result, when these heterogeneous models interact, the realized outcome can deviate substantially from the outcome each agent expects based on its own local model. In this work, we introduce the game-to-real gap, a new metric that quantifies the impact of such model misspecification in multi-agent environments. The game-to-real gap is defined as the difference between the utility an agent actually obtains in the multi-agent environment (where other agents may have misspecified models) and the utility it expects under its own game model. Focusing on quadratic network games, we show that misspecifications in either (i) the external shock or (ii) the player interaction network can lead to arbitrarily large game-to-real gaps. We further develop novel network centrality measures that allow exact evaluation of this gap in quadratic network games. Our analysis reveals that standard network centrality measures fail to capture the effects of model misspecification, underscoring the need for new structural metrics that account for this limitation. Finally, through illustrative numerical experiments, we show that existing centrality measures in network games may provide a counterintuitive understanding of the impact of model misspecification.**

I. Introduction

The efficacy of decision-making and control algorithms within multi-agent settings is conditioned on the intentions,

in response to one another [8]–and game-theoretic planning– in which an agent computes an equilibrium based on its conjectured game model, which is then used to computes its strategy [9]. In this work, we focus on the latter framework to capture decisions made by engineers and autonomous agents with significant lead time but with little opportunity to revise after deployment, e.g., an autonomous race-car which must develop a defending and overtaking policy in advance while preparing for race day event [10],or distributed generator and storage facilities participating in smart grid demand management programs [11]. Traditional game theory depends on fine assumptions about other players and can be sensitive to changes. It is our intention to provide guarantees on the robustness of game-theoretical solutions to enable engineers to more confidently make the leap from theory to use in reality. To this end, we seek to develop formal analysis methodologies that will aid in promoting design techniques within multi-agent systems that are robust to mischaracterizations of other agents' intentions or capabilities.

We formalize the idea of game-theoretic planning by assigning each agent a predictive model (consisting of a game and a solution concept) with which they may leverage optimization techniques to devise their control policy. We are interested in the case where the predictions provided by these models are inaccurate and heterogeneous among

Joint work with…

Chinmay Maheshwari
Assistant Professor
Johns Hopkins University

Manxi Wu
Assistant Professor
UC Berkeley

Shankar Sastry
Professor
UC Berkeley

Network Games with Misspecification

$$u_i \in \mathbb{R}^{m_i}, \ \forall i \in N$$

$$J_i^{(j)}(u) = \frac{1}{2} u_i^\top u_i - u_i^\top \left(P_{i,-}^{(j)} u + \epsilon_i^{(j)} \right)$$

$u^{(j)} \in \texttt{Soln}(G^{(j)})$ if it is a Nash equilibrium, i.e.,
$J_i^{(j)}(u) \leq J_i^{(j)}(u_i', u_{-i}), \ \forall u_i \in \mathcal{U}^{(j)}, \ i \in N$

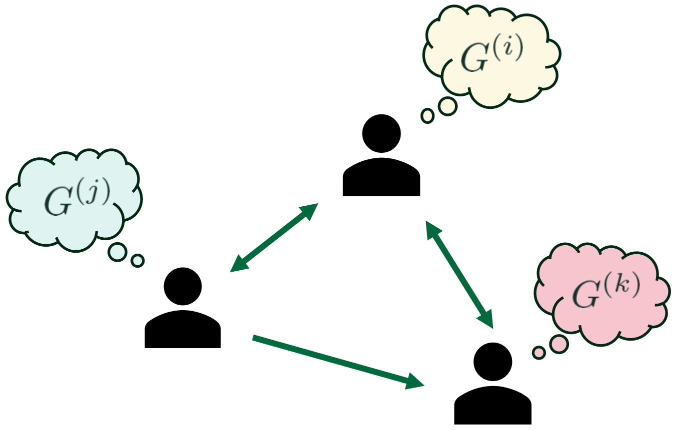

Network Games with Misspecification

$$u_i \in \mathbb{R}^{m_i}, \ \forall i \in N$$

Interaction — Exogeneous shock

$$J_i^{(j)}(u) = \frac{1}{2} u_i^\top u_i - u_i^\top \left(P_{i,-}^{(j)} u + \epsilon_i^{(j)} \right)$$

$u^{(j)} \in \texttt{Soln}(G^{(j)})$ if it is a Nash equilibrium, i.e.,
$J_i^{(j)}(u) \leq J_i^{(j)}(u_i', u_{-i}), \ \forall u_i \in \mathcal{U}^{(j)}, \ i \in N$

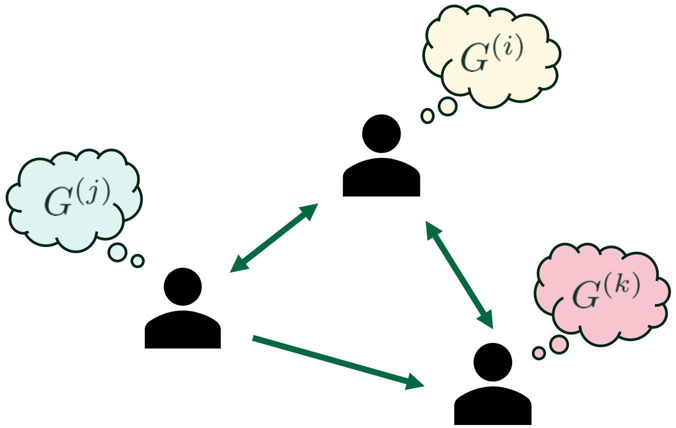

Forecast Misspecification

$$\epsilon^{(i)} \neq \epsilon^{(j)}$$

Interaction Misspecification

$$P^{(i)} \neq P^{(j)}$$

Network Games with Misspecification

Proposition 1: In quadratic network aggregative games, for any $\delta, M > 0$, there exits some $\{\epsilon^{(i)}\}_{i \in N}$ and interaction matrix P such that

$$\|\epsilon^{(i)} - \epsilon^{(j)}\|_2 \leq \delta \quad \forall\ i, j \in N,$$

but that for each $i \in N$,

$$J_i(u^{\circ}) - J_i(u^{(i)}) > M.$$

Network Games with Misspecification

Proposition 1: In quadratic network aggregative games, for any $\delta, M > 0$, there exits some $\{\epsilon^{(i)}\}_{i \in N}$ and interaction matrix P such that

$$\|\epsilon^{(i)} - \epsilon^{(j)}\|_2 \leq \delta \quad \forall\, i, j \in N,$$

but that for each $i \in N$,

$$J_i(u^\circ) - J_i(u^{(i)}) > M.$$

$$P = \begin{bmatrix} 0 & 0 & \gamma \\ \gamma & 0 & 0 \\ 0 & \gamma & 0 \end{bmatrix},$$

$$\epsilon^{(1)} = \begin{bmatrix} 1 \\ \beta \\ \beta \end{bmatrix}, \ \epsilon^{(2)} = \begin{bmatrix} \beta \\ 1 \\ \beta \end{bmatrix}, \ \epsilon^{(3)} = \begin{bmatrix} \beta \\ \beta \\ 1 \end{bmatrix}$$

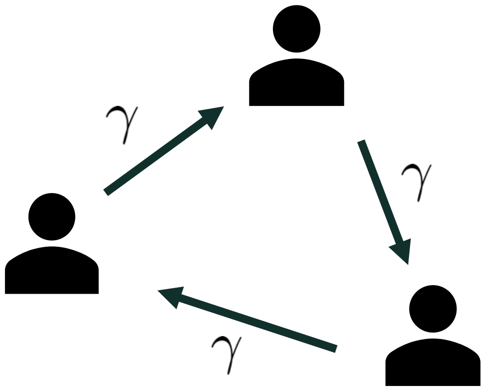

$$\lim_{\gamma \nearrow 1} J_i(u^\circ) - J_i(u^{(i)}) = \infty$$

Observation: embedded cycles cause a larger game-to-real gap

Network Games with Misspecification

Proposition 1: In quadratic network aggregative games, for any $\delta, M > 0$, there exits some $\{\epsilon^{(i)}\}_{i \in N}$ and interaction matrix P such that

$$\|\epsilon^{(i)} - \epsilon^{(j)}\|_2 \leq \delta \quad \forall\, i, j \in N,$$

but that for each $i \in N$,

$$J_i(u^\circ) - J_i(u^{(i)}) > M.$$

Definition 1: The *Shock Misspecification Centrality* between player i and j is denoted by

$$\mathcal{B}_{i,j} = [(I-P)^{-1}]_{i,-}^\top P_{i,j} [(I-P)^{-1}]_{j,-} \in \mathbb{R}^{m \times m}.$$

Theorem 2 (abrv.): The game-to-real gap is characterized by

$$J_i(u^\circ) - J_i(u^{(i)}) = \sum_{j \neq i} \epsilon^{(i)\top} \mathcal{B}_{i,j} (\epsilon^{(i)} - \epsilon^{(j)}).$$

Graphs centrality measures are insufficient for characterizing game-to-real gap, need pair-wise centrality

Network Games with Misspecification

Proposition 1: In quadratic network aggregative games, for any $\delta, M > 0$, there exits some $\{\epsilon^{(i)}\}_{i \in N}$ and interaction matrix P such that

$$\|\epsilon^{(i)} - \epsilon^{(j)}\|_2 \leq \delta \quad \forall\, i, j \in N,$$

but that for each $i \in N$,

$$J_i(u^\circ) - J_i(u^{(i)}) > M.$$

Definition 1: The *Shock Misspecification Centrality* between player i and j is denoted by

$$\mathcal{B}_{i,j} = [(I-P)^{-1}]_{i,-}^\top P_{i,j} [(I-P)^{-1}]_{j,-} \in \mathbb{R}^{m \times m}.$$

Theorem 2 (abrv.): The game-to-real gap is characterized by

$$J_i(u^\circ) - J_i(u^{(i)}) = \sum_{j \neq i} \epsilon^{(i)\top} \mathcal{B}_{i,j} (\epsilon^{(i)} - \epsilon^{(j)}).$$

Graphs centrality measures are insufficient for characterizing game-to-real gap, need pair-wise centrality

Interaction Misspecification:

- Proposition 3: arbitrarily bad game-to-real gap for any bounded norm of graph difference
- Proposition 4: characterization of game-to-real gap by weighted pair-wise centrality measure

Simultaneous Misspecification:

- Corollary 1: characterization of game-to-real gap

Ongoing work: more general game-to-real gap by parametric VI analysis, apply to LQ games, network routing games, and zero-sum games

Designing Around Misspecification

- Interaction Sample complexity

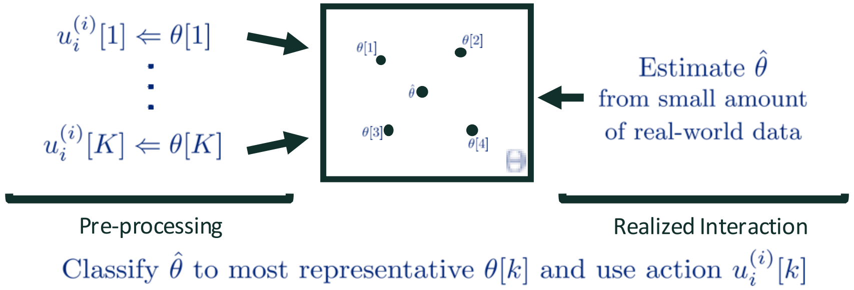

- Robust game-model selection

$$\underset{\theta^{(i)} \in \Theta}{\text{minimize}} \quad \max_{\theta^{(-i)} \sim \Theta_{-i}} \left[J_i \left(u^{\circ} \left(\theta^{(i)}, \theta^{(-i)} \right) \right) \right]$$

$$\underset{\theta^{(i)} \in \Theta}{\text{minimize}} \quad \mathbb{E}_{\theta^{(-i)} \sim \mu_i} \left[J_i \left(u^{\circ} \left(\theta^{(i)}, \theta^{(-i)} \right) \right) \right]$$

- Online planning
  - Model predictive games

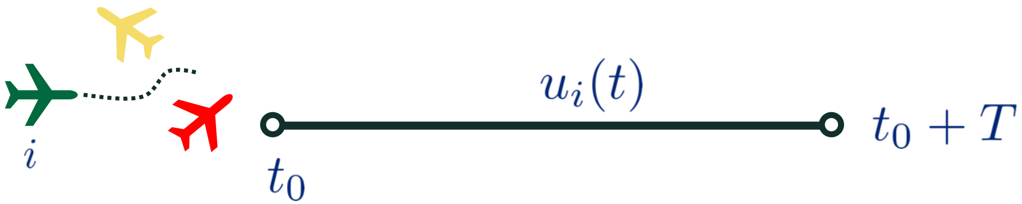

Ada Yildirim
PhD Student
Dartmouth College

Model Predictive Control (MPC)

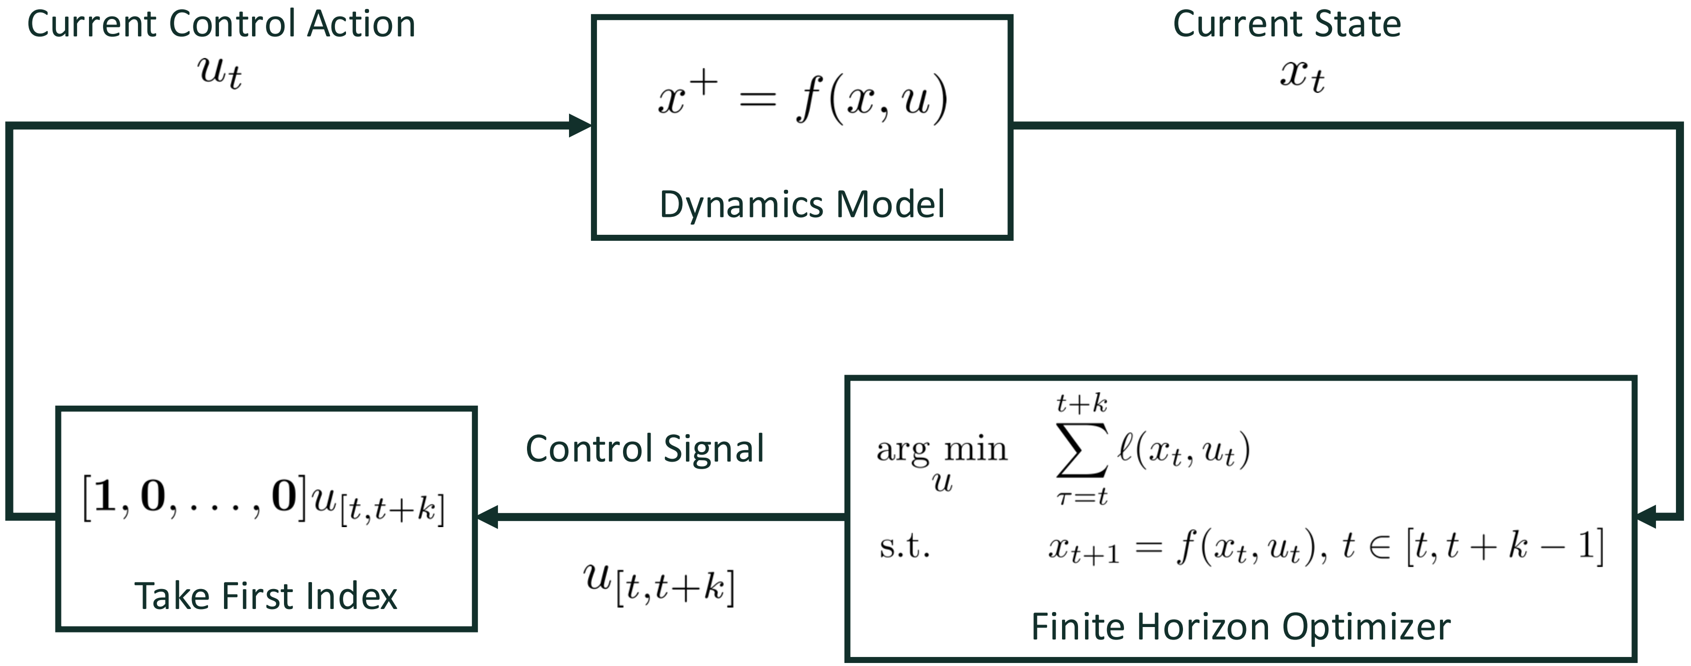

Model Predictive Control (MPC)

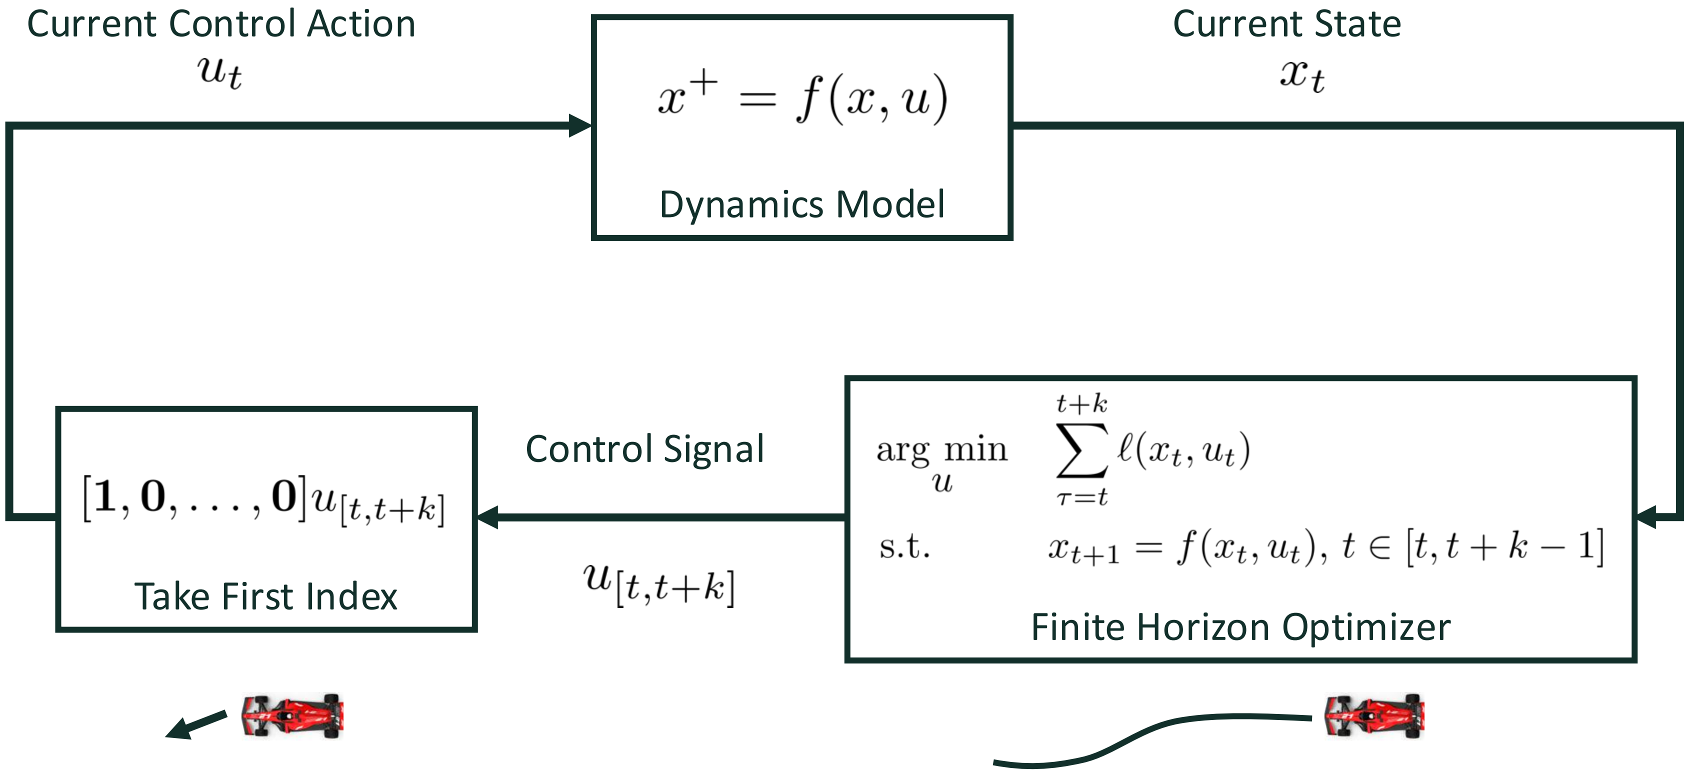

Solve for optimal control signal over finite horizon and deploy first timestep

Game-Theoretic MPC

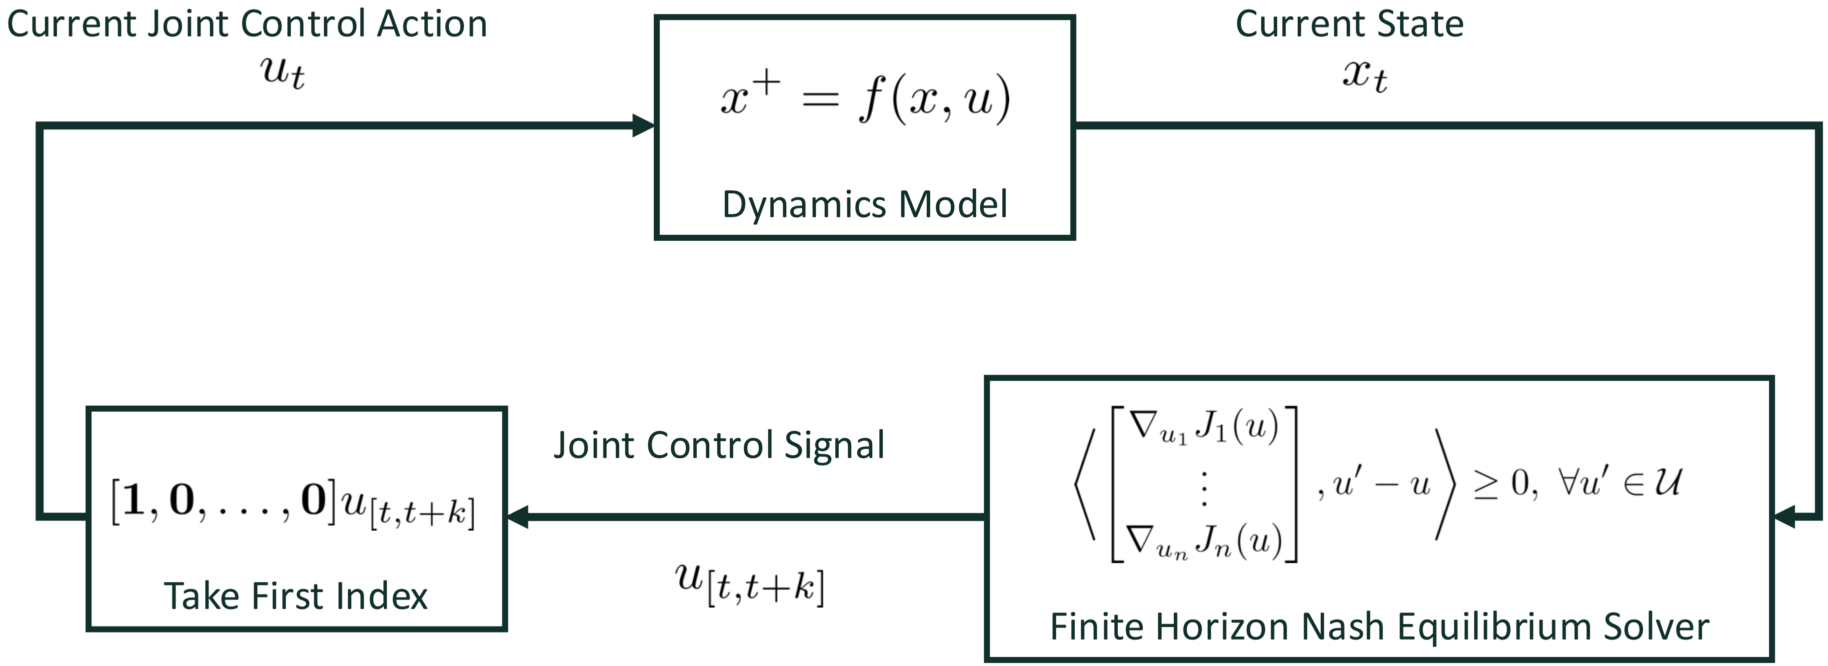

Hall, Sophie, Dominic Liao-McPherson, Giuseppe Belgioioso, and Florian Dörfler. "Stability Certificates for Receding Horizon Games." *arXiv preprint arXiv:2404.12165* (2024).
E. Benenati and G. Belgioioso. "The explicit game-theoretic linear quadratic regulator for constrained multi-agent systems." arXiv preprint arXiv:2512.07749 (2025).

Game-Theoretic MPC

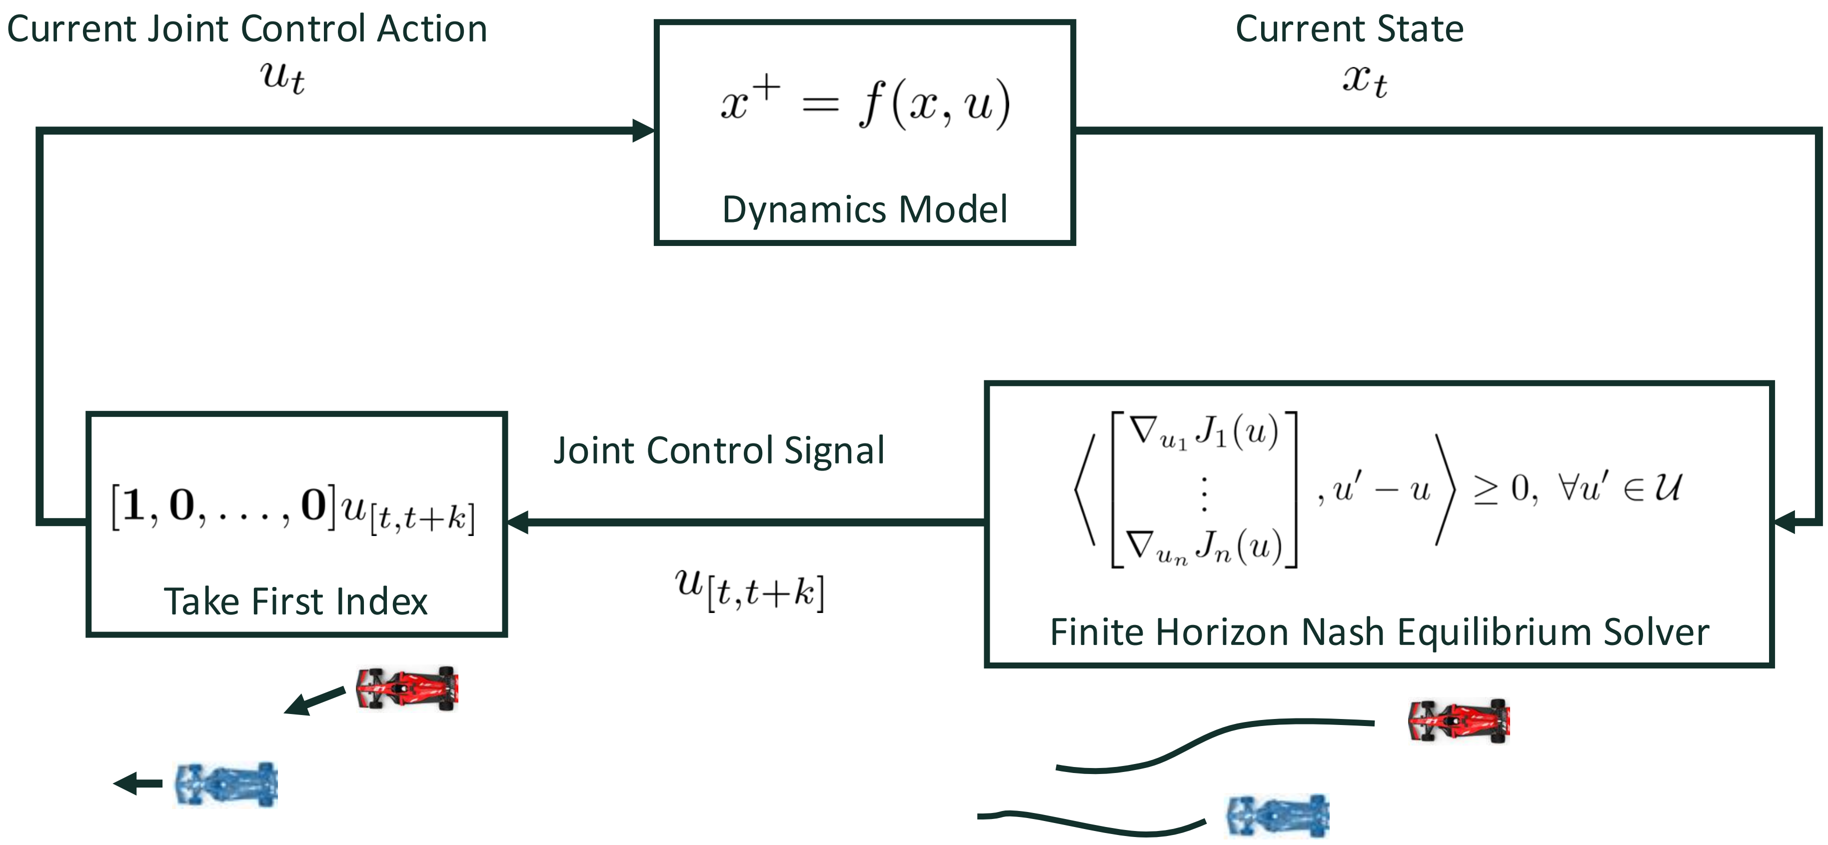

Solve for ***Nash Equilibrium*** control signal over finite horizon and deploy first timestep

Hall, Sophie, Dominic Liao-McPherson, Giuseppe Belgioioso, and Florian Dörfler. "Stability Certificates for Receding Horizon Games." *arXiv preprint arXiv:2404.12165* (2024).
E. Benenati and G. Belgioioso. "The explicit game-theoretic linear quadratic regulator for constrained multi-agent systems." arXiv preprint arXiv:2512.07749 (2025).

Misspecificed Game-Theoretic MPC

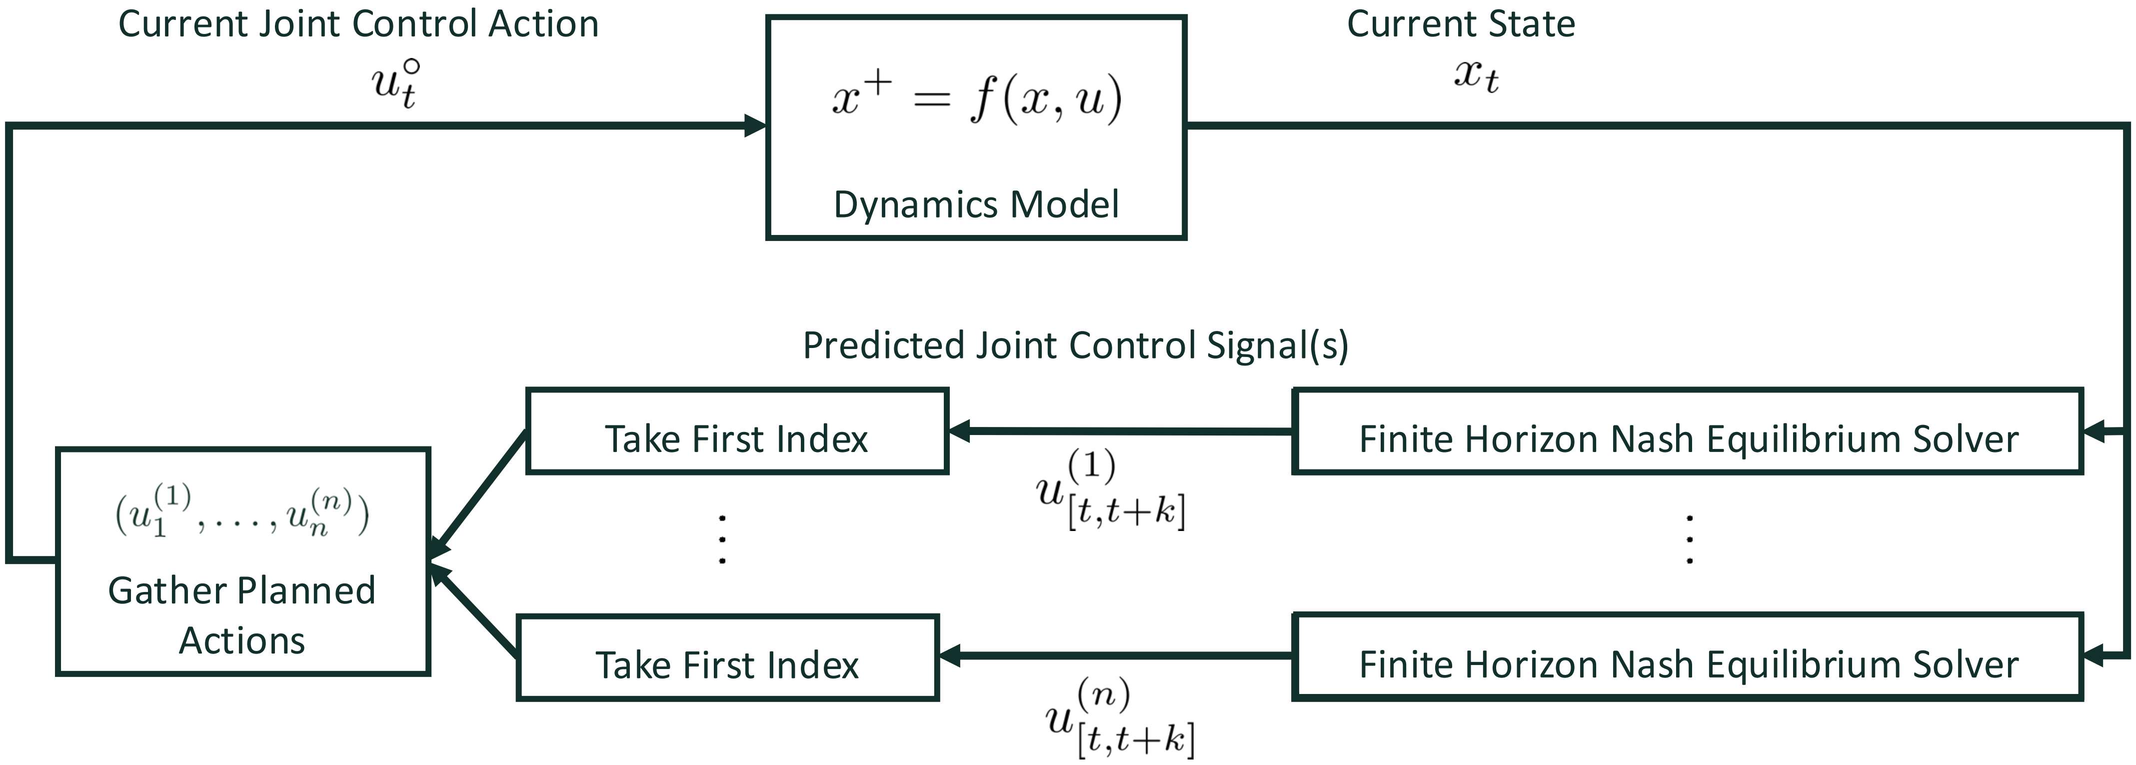

Yildirim, Ada, and Bryce L. Ferguson. "Stability and Sensitivity Analysis for Objective Misspecifications Among Model Predictive Game Controllers." *arXiv preprint arXiv:2604.08303*(2026).

Misspecified Game-Theoretic MPC

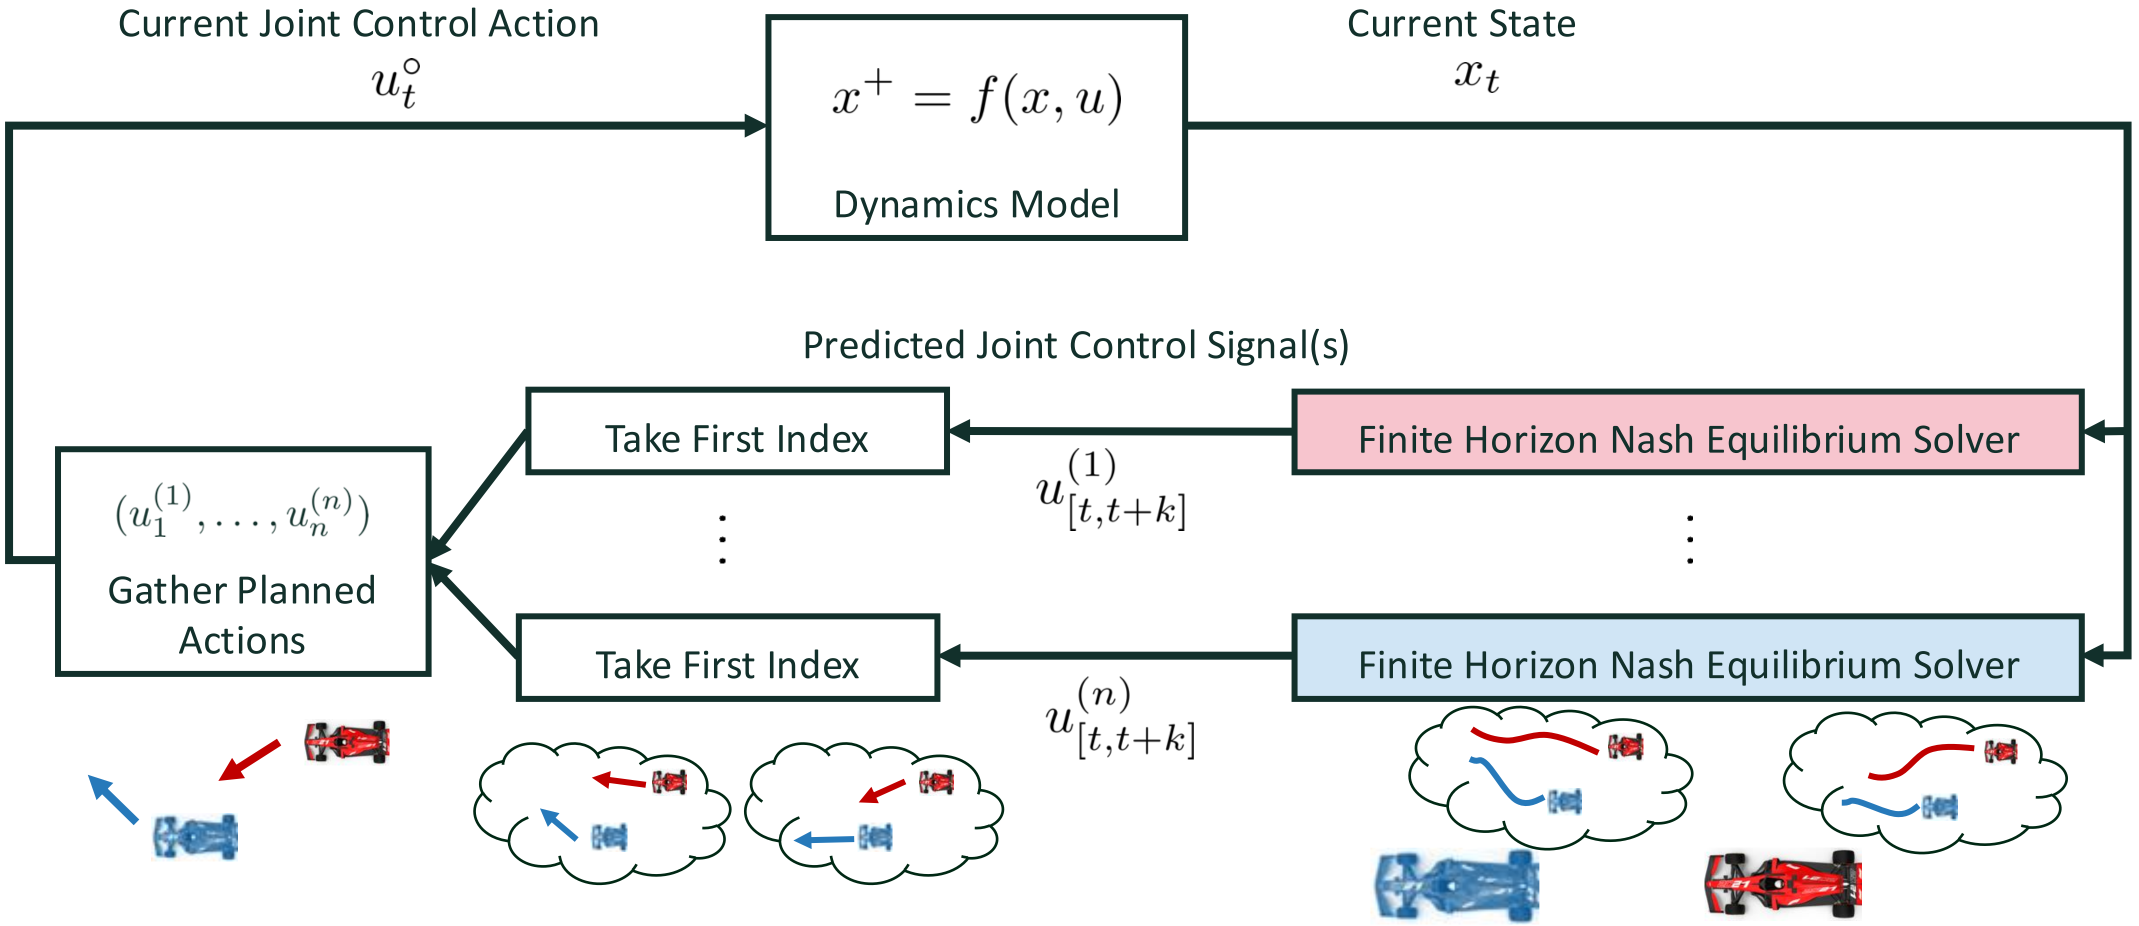

Each player solves for a ***Nash Equilibrium*** control signal over finite horizon and deploys ***respective*** first timestep

Stability Criteria with Misspecifications

Theorem: Under monotonicity assumptions, and there exist a positive-definite matrix $P \succ 0$ and a scalar $\lambda > 0$, such that

$$\begin{bmatrix} A^\top PA - P & A^\top P\hat{B} \\ \hat{B}^\top PA & \hat{B}^\top P\hat{B} \end{bmatrix} + \lambda W \preceq -\varepsilon I,$$

for some $\varepsilon > 0$, where ρ_i is the strong monotonicity constant of $F_u^{(j)}$, $\hat{B} = \begin{bmatrix} B_1\Xi_1 & B_2\Xi_2 & \dots & B_n\Xi_n \end{bmatrix}$, and the blocks of W are defined as, $\forall i \in N$,

$$W_{1(j+1)} = -\frac{F_x^{(j)}}{2}, \; W_{(j+1)1} = -\frac{(F_x^{(j)})^T}{2},$$
$$W_{(j+1)(j+1)} = -\rho_j I,$$

and zero otherwise, where the subscript indices denote the block positions for W, then there exists a globally asymptotically stable equilibrium point $\bar{x} \in \mathbb{R}^{n_x}$

Stability criteria depends on a non-trivial combination of each player's game

Dynamics with Misspecifications

Games with little misspecification

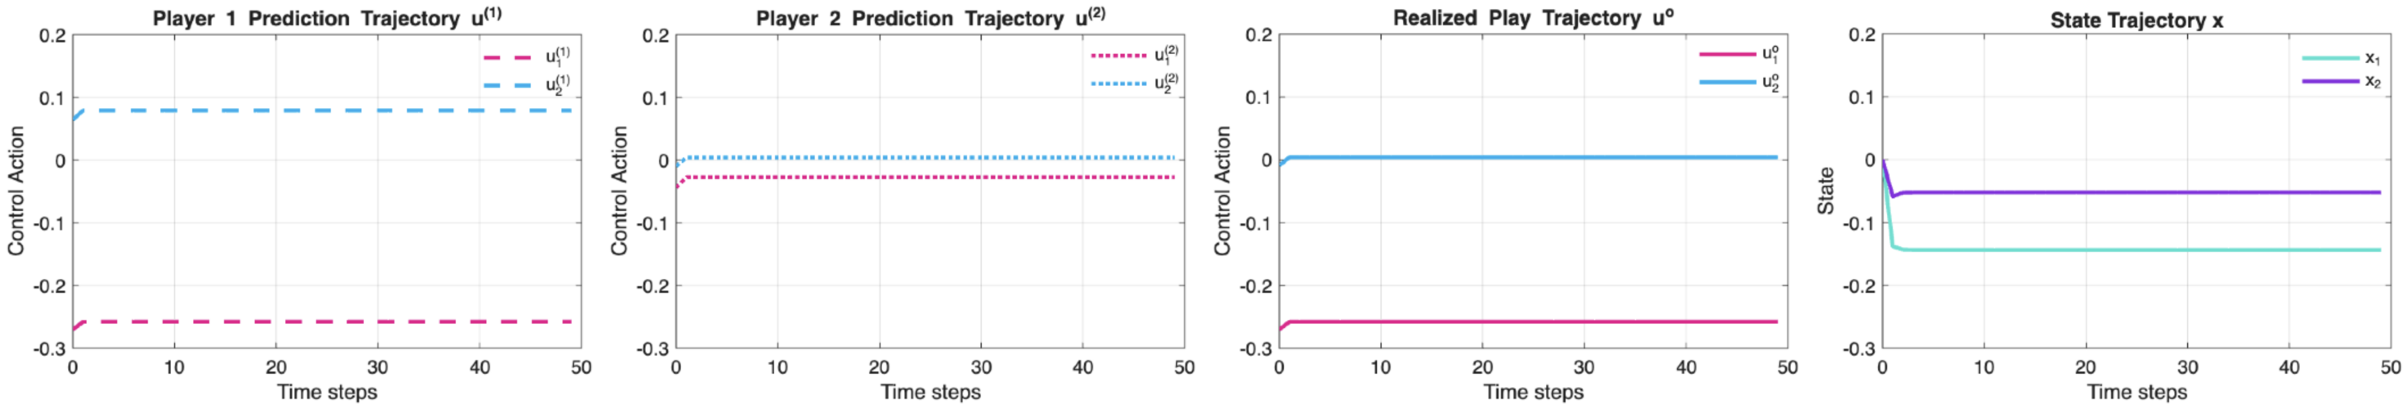

Games with large misspecifications

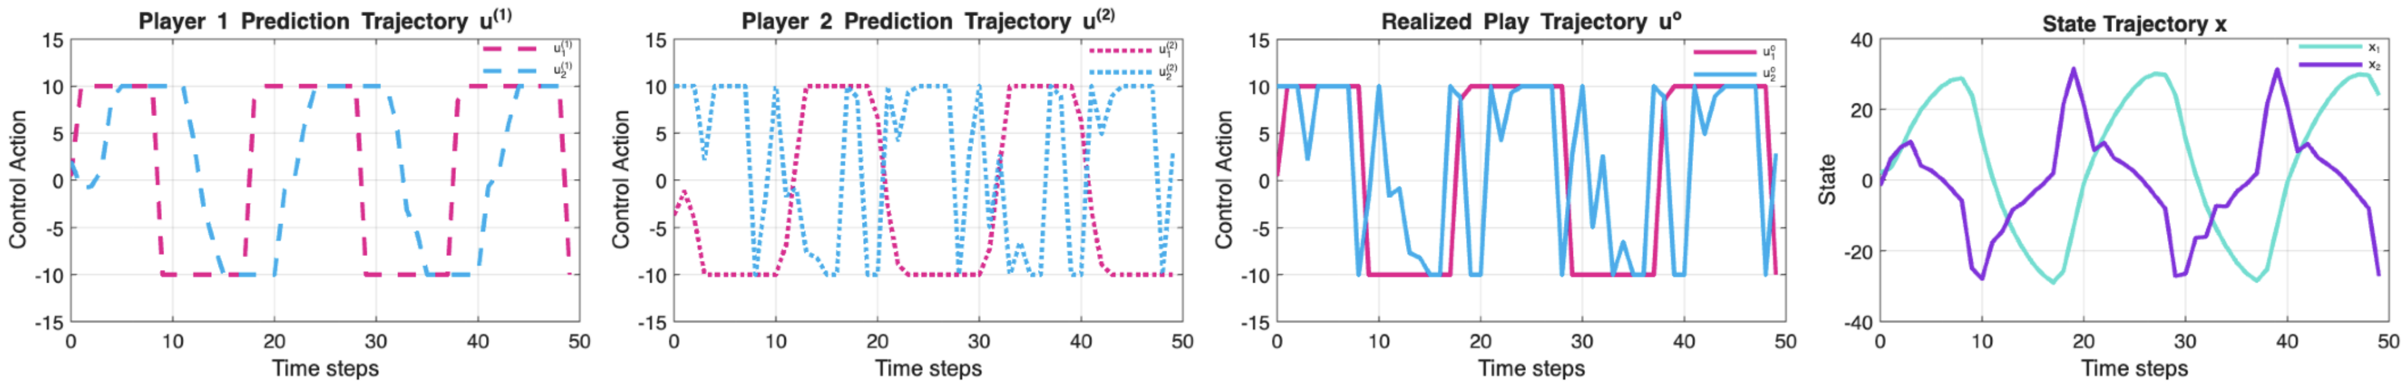

Sensitivity of Equilibrium to Misspecifications

$G(x_0, \theta^{(j)})$

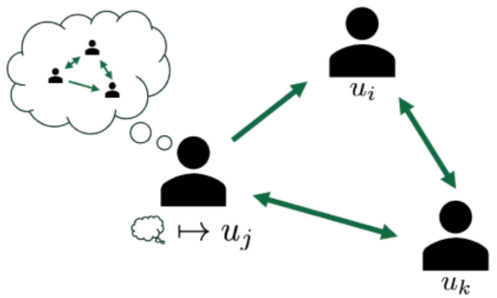

$$g_{i,t}^{(j)}(x_t, u_t) = \sum_{k \in |\theta^{(j)}|} \theta_k^{(j)} \left(x_t^\top Q_i^k x_t + x_t^\top q_i^k + u_t^\top R_i^k u_t \right)$$

Proposition: For parameter $\bar{\theta} = (\theta^{(i)})_{i \in N}$, let $x^\star(\bar{\theta})$ be a unique equilibrium of the heterogeneous MPG system with conjectured games $\{G(\cdot, \theta^{(i)})\}_{i \in N}$. Under mild assumptions, then the sensitivity of $x^\star$ to δ is:

$$\nabla_{\bar{\theta}} x^*(\bar{\theta}) = \left(I - TB\,\Xi\,\nabla_x \bar{u}(\bar{\theta}, x) \right)^{-1} T\,\Xi\,\nabla_{\bar{\theta}} \bar{u}(\bar{\theta}, x)$$

where $T := (I - A)^{-1} B$.

Equilibria with Misspecifications

How does the equilibrium of misspecified GT MPC change with greater misalignment?

$$G^{(1)} = G^A \qquad G^{(2)} = \theta G^A + (1-\theta)G^B$$

No misspecification

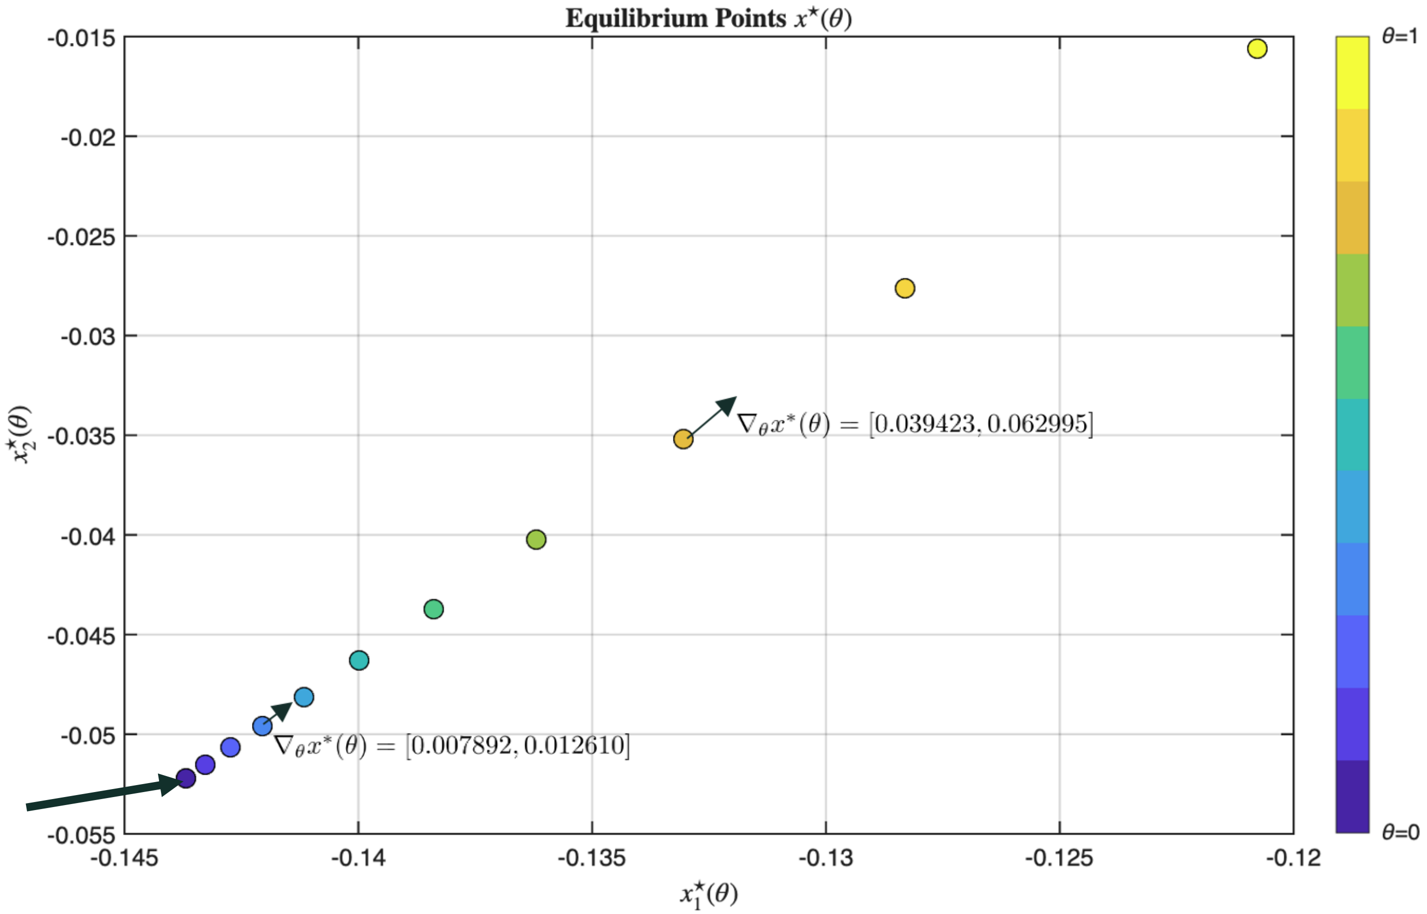

Conclusion

- Collaborative architectures between centralized and distributed systems can improve system performance
- Competitive effectiveness is sensitive to misspecification in the prediction model

Future work:

- Dynamically learning collaboration architecture
- Inverse learning multi-agent prediction models from interaction data
- Joint competitive-cooperative multi-agent environments

Bryce.L.Ferguson@dartmouth.edu