Intelligent Coordination of Heterogeneous Sensors in Advanced Power Systems

PI: Kagan Tumer

Oregon State University

kagan.tumer@oregonstate.edu

Agreement Number: DE-FE0011403

NETL Project Manager: Maria Reidpath

Motivation

- Where are we?
  - Advanced energy systems becoming more interconnected
    - Larger, distributed, more stochastic
  - Computation pushed further down the pipe
    - More powerful, cheaper, smaller devices

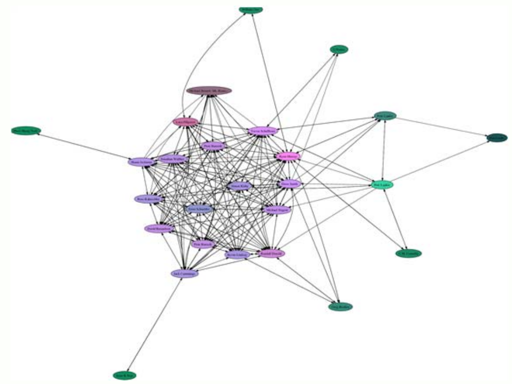

Motivation

- Where are we?
  - Advanced energy systems becoming more interconnected
    - Larger, distributed, more stochastic
  - Computation pushed further down the pipe
    - More powerful, cheaper, smaller devices
- Where are we going?
  - Hybrid systems
    - Electrical/bio/mechanical devices
  - Smart sensors
    - Tens of thousands of tiny, simple, unreliable sensors

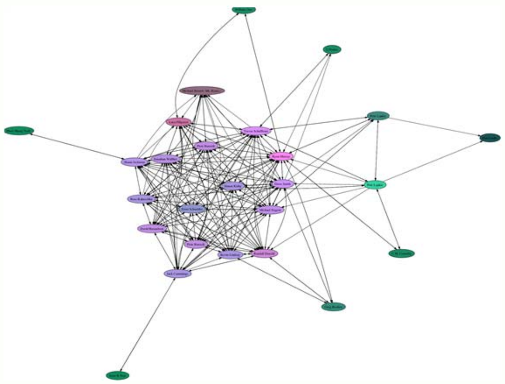

Motivation

- Where are we?
  - Advanced energy systems becoming more interconnected
    - Larger, distributed, more stochastic
  - Computation pushed further down the pipe
    - More powerful, cheaper, smaller devices
- Where are we going?
  - Hybrid systems
    - Electrical/bio/mechanical devices
  - Smart sensors
    - Tens of thousands of tiny, simple, unreliable sensors
- What do we need to account for?
  - Tens of thousands of sensors
  - Failing sensors
  - Dynamic and stochastic environments

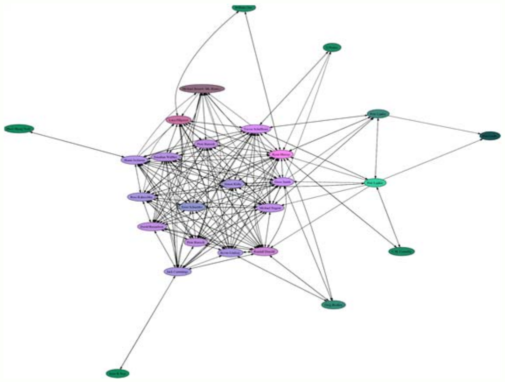

Key Challenge

How do we coordinate a very large number
of heterogeneous sensors and actuators so that they
collectively optimize a system objective function?

Where Should Focus Be?

- New optimization algorithms?

- New control algorithms?

Where Should Focus Be?

- New optimization algorithms?

No!

- New control algorithms?

No!

Where Should Focus Be?

- New optimization algorithms?

No!

- New control algorithms?

No!

- Focus on:
  - How to control?
  - What to optimize?
  - What are “good system” properties?

Cooperative Multiagent Systems

- System Description:
  - Each sensor has an ***agent objective*** it aims to optimize
  - A ***system objective*** rates the entire system's performance

- Important issues:
  - *How do we set agent objective functions?*
  - *How to update them?*
  - *Can agents compute those objective functions?*
  - *What happens when information is missing?*
  - *What happens when agents fail?*
  - *What happens when system goals change?*

Outline

- Motivation
- Critical Concepts
- Project Objectives
- Objective 1: Methodology and Results
- Objective 2: Methodology and Results
- Closing Remarks

Critical Concepts

- ***Evolutionary Algorithms***
- ***Cooperative Coevolutionary Algorithms***
- Multiagent Reinforcement Learning
- Objectives in Self-Organizing Systems
- Difference Evaluation Functions

Evolutionary Algorithms

- Stochastic, population-based search algorithm
- Operators: Mutation, Fitness Assignment, Selection
- Work well in optimization problems where gradient information is unavailable

Evolutionary Algorithms

Initialize *n* policies

Evolutionary Algorithms

Initialize *n* policies

Mutation

Create *n* slightly altered policies (2*n*)

Evolutionary Algorithms

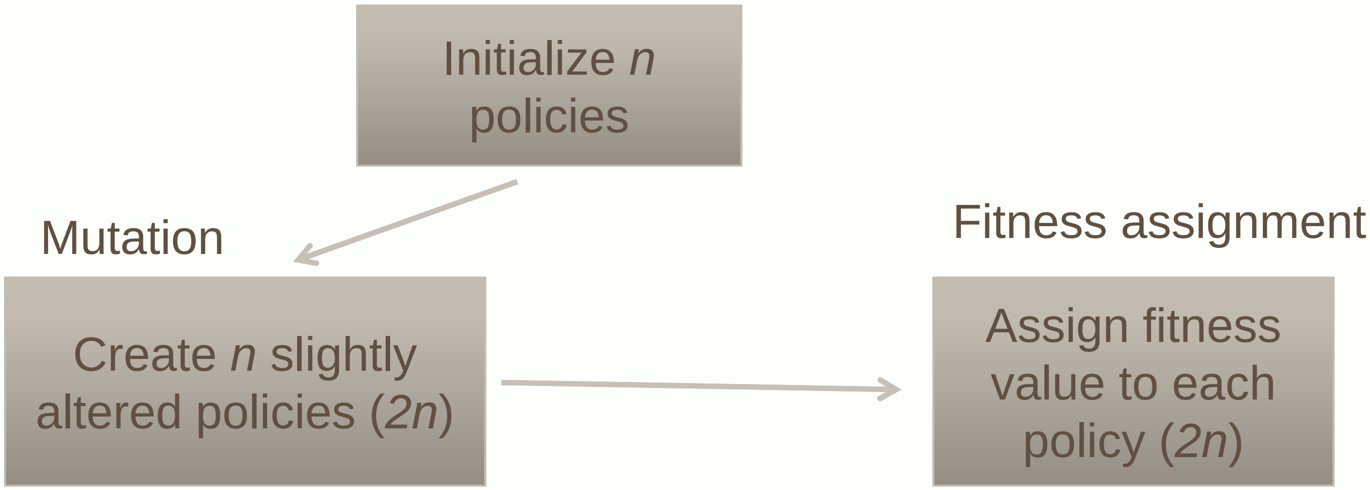

Evolutionary Algorithms

Initialize *n* policies

Mutation

Create *n* slightly altered policies (2*n*)

Fitness assignment

Assign fitness value to each policy (2*n*)

Selection

Select *n* policies to survive to next generation (*n*)

Evolutionary Algorithms

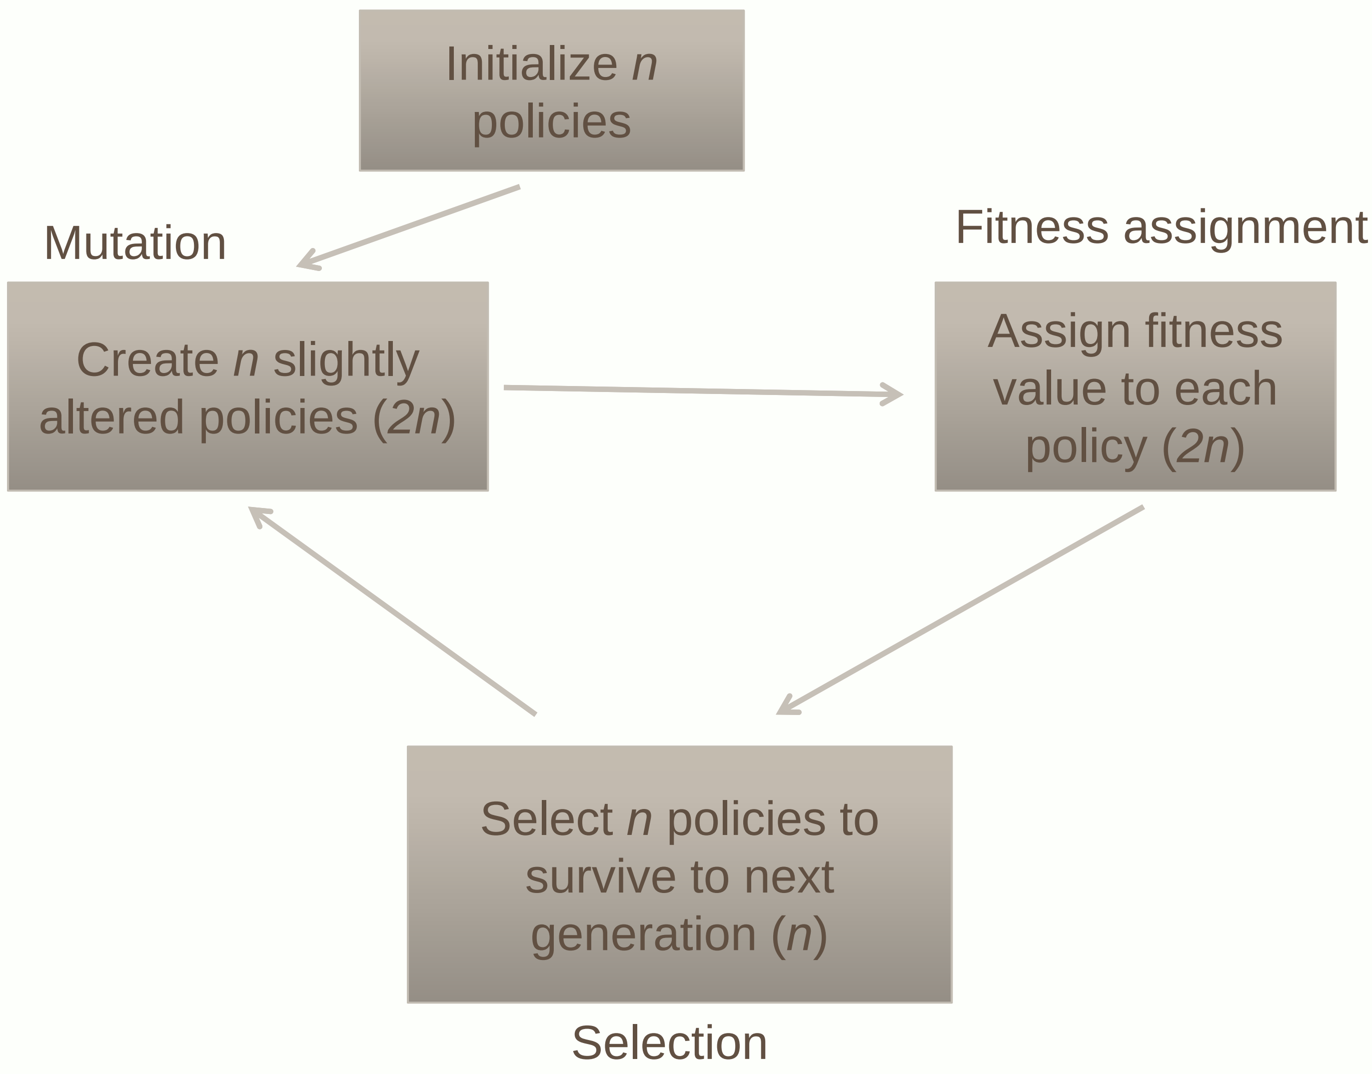

Cooperative Coevolutionary Algorithms

- Evolutionary algorithms need to be extended for many agents interacting
- Multiple coupled evolutionary algorithms in parallel
- Only significant difference from standard evolutionary algorithm is fitness assignment stage

Cooperative Coevolutionary Algorithms

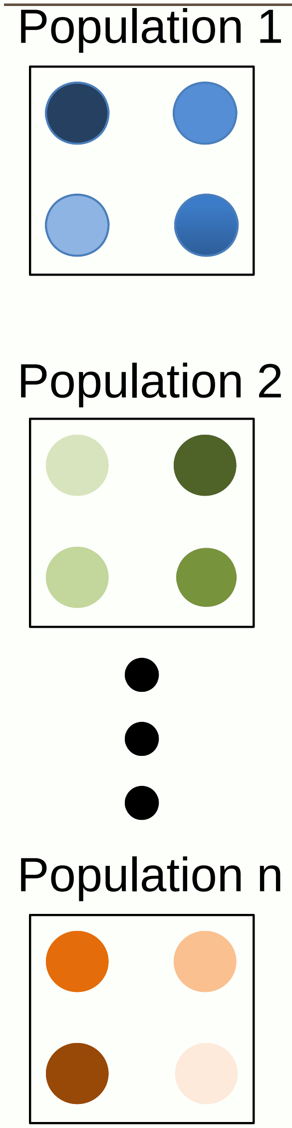

Cooperative Coevolutionary Algorithms

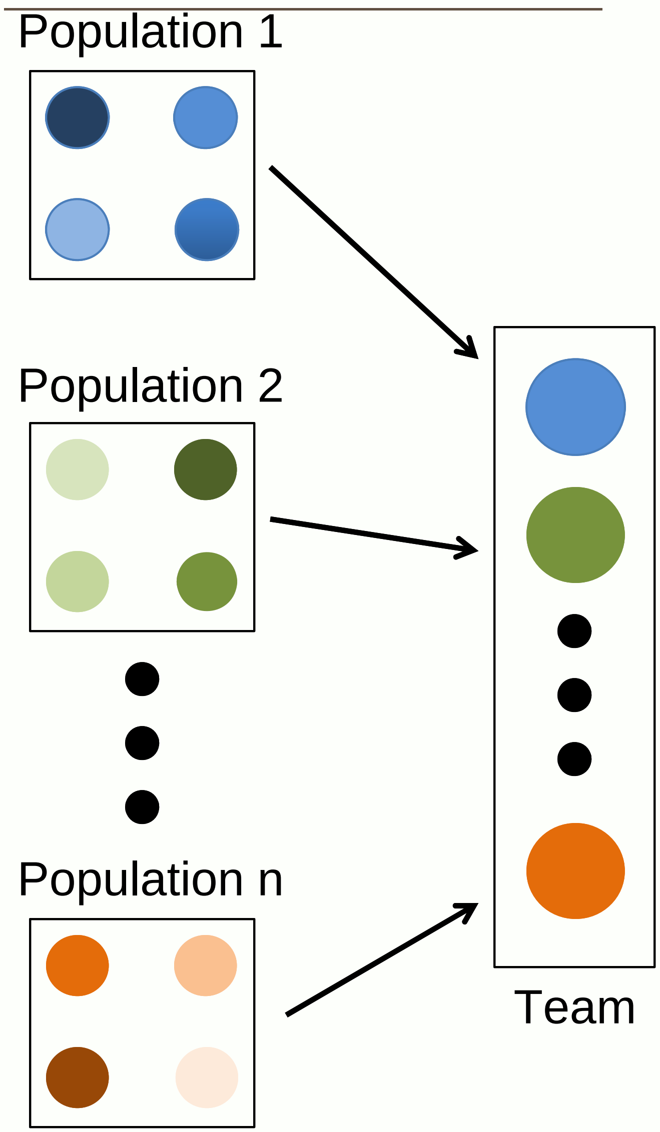

Cooperative Coevolutionary Algorithms

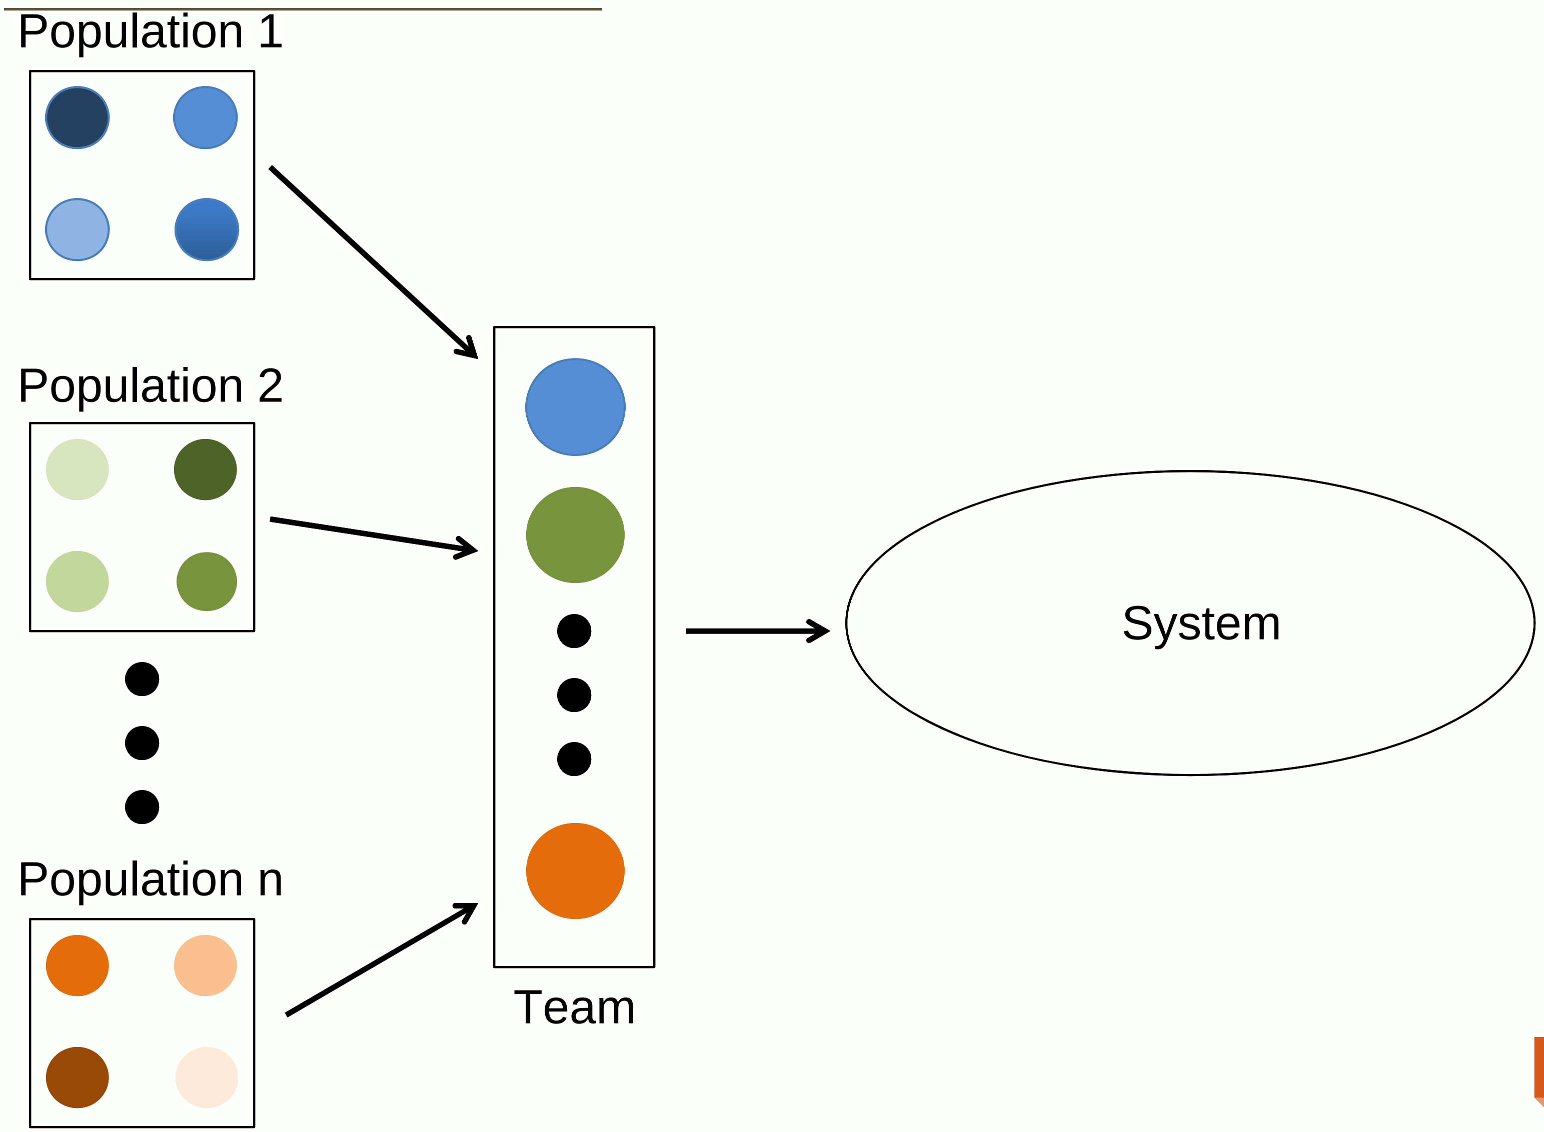

Cooperative Coevolutionary Algorithms

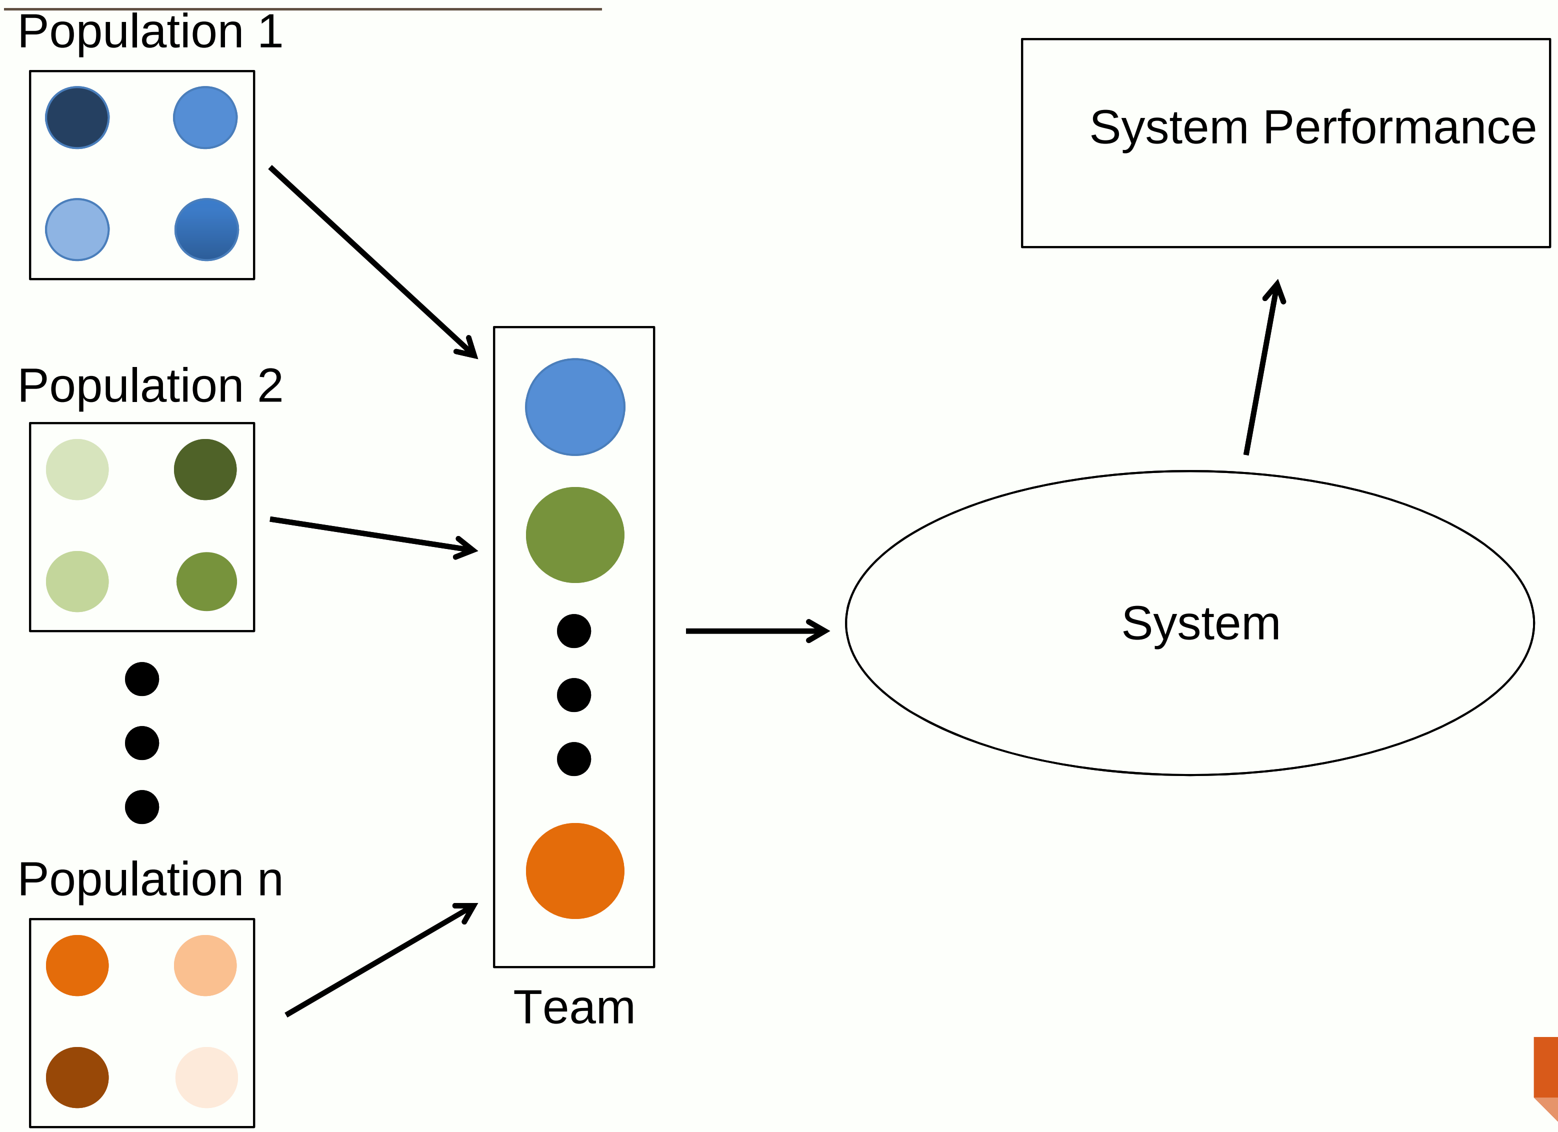

Cooperative Coevolutionary Algorithms

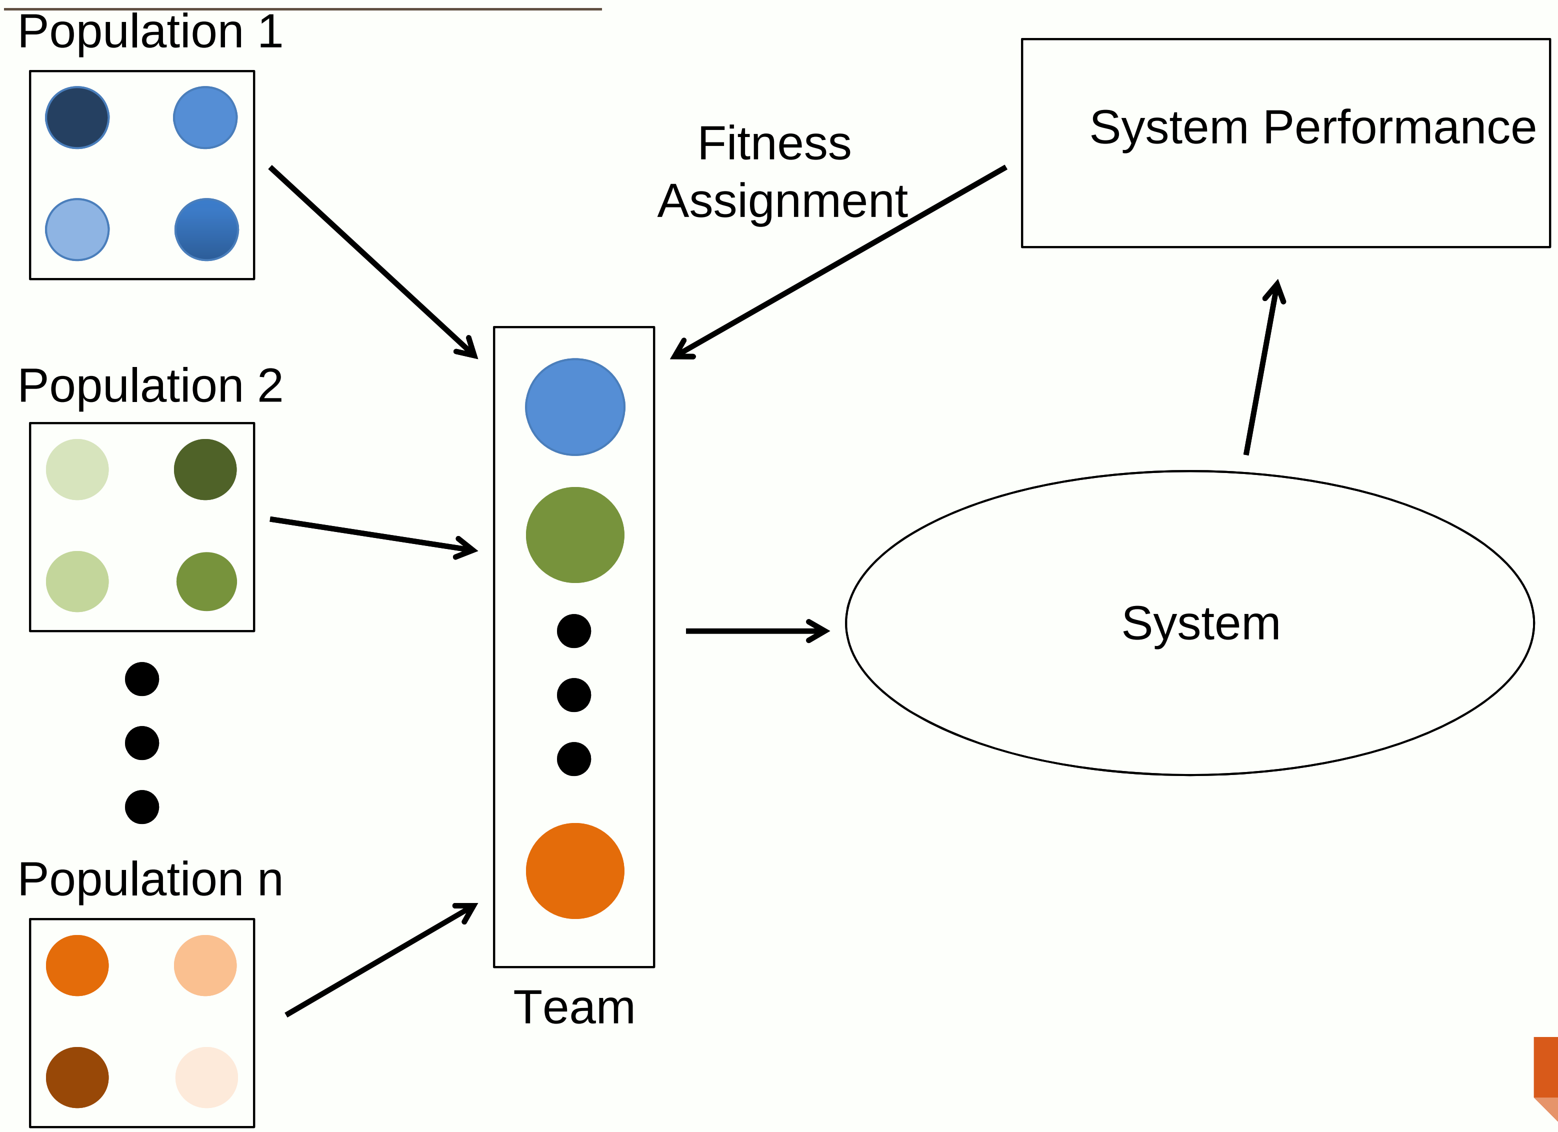

Cooperative Coevolutionary Algorithms

- Fitness of an agent is a function of two things:
  - The agent's policy
  - How the collaborating agents act
- Fitness assignment in cooperative coevolutionary algorithms is very context-dependent and subjective
- Credit assignment problem extremely difficult to solve
  - Fitness function shaping

Critical Concepts

- Evolutionary Algorithms
- Cooperative Coevolutionary Algorithms
- **Multiagent Reinforcement Learning**
- Objectives in Self-Organizing Systems
- Difference Evaluation Functions

Multiagent Reinforcement Learning

- A set of autonomous agents learns to coordinate/self-organize
- Model-free method to develop controllers for distributed systems
- Agents conduct trials repeatedly, and learn which actions yield high performance

Multiagent Reinforcement Learning

- Multiagent Reinforcement Learning:
  - Each agent maintains a Q-table: maps actions to their expected utility
  - After taking an action and receiving feedback, update Q-table:

$$Q(a) \leftarrow \alpha R + (1 - \alpha) Q(a)$$

- Key problems in multiagent learning:
  - Need to ensure agents don't work at cross-purposes
  - Need to ensure each agent contributes to the system
  - Setting agent **objectives** is a nontrivial task, and choice of objective functions has a large impact on system performance

Example: Global Objective Function

- Each agent receives the overall system performance as feedback
- Problem: too much noise in the feedback signal

- A team of 100 agents is acting in an environment:
  - 99 agents act optimally
  - 1 agent does nothing
  - Overall, the system performs well, and the agent that did nothing believes it helped the system

Example: Local Objective Function

- Each agent receives feedback based on local performance measures
- Problem: agents can become "greedy," and act to harm the system

- Agents acting in a surveillance domain
  - Local feedback based on "how much" information an agent collects
  - Agents will learn to fight over the easy to observe measurements, rather than distributing their efforts across the system

Objectives in Self-Organizing Systems

- Multiagent Learning
  - Each agent has a local objective it needs to optimize
  - Coevolutionary algorithms: fitness function
  - Reinforcement learning: reward signal

- We have seen that improper choice of fitness/reward can lead to poor system performance
  - Global feedback: too noisy
  - Local feedback: can lead to agents working at cross-purposes

- What should to agent feedback be?

Desirable Objective Function Properties

- $g_i(z)$ should be *aligned* with $G(z)$
  - An agent which increases $g_i(z)$ also increases $G(z)$
  - "Is what's good for me good for the full system?"

Desirable Objective Function Properties

- $g_i(z)$ should be *aligned* with $G(z)$
  - An agent which increases $g_i(z)$ also increases $G(z)$
  - "Is what's good for me good for the full system?"

$$\boxed{\mathcal{F}_{g_i} = \frac{\sum_{z'} u[((g_i(z) - g_i(z'))(G(z) - G(z'))]}{\sum_{z'} 1}}$$

Desirable Objective Function Properties

- $g_i(z)$ should be *aligned* with G(z)
  - An agent which increases $g_i(z)$ also increases G(z)
  - "Is what's good for me good for the full system?"

$$\boxed{\mathcal{F}_{g_i} = \frac{\sum_{z'} u[((g_i(z) - g_i(z'))(G(z) - G(z'))]}{\sum_{z'} 1}}$$

- $g_i(z)$ should be *sensitive* to agent's actions
  - Signal to noise, locality
  - "Can I extract what's good for me from signal?"

Desirable Objective Function Properties

- $g_i(z)$ should be *aligned* with $G(z)$
  - An agent which increases $g_i(z)$ also increases $G(z)$
  - "Is what's good for me good for the full system?"

$$\mathcal{F}_{g_i} = \frac{\sum_{z'} u[((g_i(z) - g_i(z'))(G(z) - G(z'))]}{\sum_{z'} 1}$$

- $g_i(z)$ should be *sensitive* to agent's actions
  - Signal to noise, locality
  - "Can I extract what's good for me from signal?"

$$L(g_i, z, z') = \frac{\|g_i(z) - g_i(z - z_i + z'_i)\|}{\|g_i(z) - g_i(z' - z'_i + z_i)\|}$$

$$L(g_i, z) = \frac{\sum_{z'} L(g_i, z, z')}{\sum_{z'} 1}$$

Difference Evaluation Functions

- Difference evaluation function defined as:

$$g_i(\mathrm{z}) = G(\mathrm{z}) - G(\mathrm{z}_{-i} + c_i)$$

Difference Evaluation Functions

- Difference evaluation function defined as:

$$g_i(\text{z}) = G(\text{z}) - G(\text{z}_{-i} + c_i)$$

- $g_i(z)$ removes portions of $G(z)$ not related to agent *i*
  - Good signal to noise ratio → $g_i(z)$ is ***sensitive*** to agent *i*'s actions

Difference Evaluation Functions

- Difference evaluation function defined as:

$$g_i(\mathsf{z}) = G(\mathsf{z}) - G(\mathsf{z}_{-i} + c_i)$$

- $g_i(z)$ removes portions of G(z) not related to agent *i*
  - Good signal to noise ratio → $g_i(z)$ is ***sensitive*** to agent *i*'s actions

- If $g_i(z)$, G(z) are differentiable, then:

$$\frac{\partial G(\mathsf{z}_{-i} + c_i)}{\partial \mathsf{z}_i} = 0 \Rightarrow \frac{\partial g_i(\mathsf{z})}{\partial \mathsf{z}_i} = \frac{\partial G(\mathsf{z})}{\partial \mathsf{z}_i}$$

Difference Evaluation Functions

- Difference evaluation function defined as:

$$g_i(z) = G(z) - G(z_{-i} + c_i)$$

- $g_i(z)$ removes portions of G(z) not related to agent *i*
  - Good signal to noise ratio → $g_i(z)$ is ***sensitive*** to agent *i*'s actions

- If $g_i(z)$, G(z) are differentiable, then:

$$\frac{\partial G(z_{-i} + c_i)}{\partial z_i} = 0 \Rightarrow \frac{\partial g_i(z)}{\partial z_i} = \frac{\partial G(z)}{\partial z_i}$$

  - Increasing $g_i(z)$ increases G(z) → $g_i(z)$ is ***aligned*** with G(z)

Where are We Now?

- Proper objective functions significantly improve system performance
- Difference evaluation functions are extremely scalable, up to network sizes of 10,000 devices

- **What about heterogeneous sensors?**

What About Heterogeneous Sensors?

- What if we have heterogeneous sensors (agents)?
  - Different capabilities
  - (Potentially) different goals

- Example: pressure and temperature sensors
  - Set of temperature sensors and pressure sensors must be optimally located in a plant
  - Aim to maximize accuracy of temperature and pressure measurements
  - What if location for optimal pressure sensor placement corresponds to location of optimal temperature sensor placement?

What About Heterogeneous Sensors?

- What if we have heterogeneous sensors (agents)?
  - Different capabilities
  - (Potentially) different goals

- Example: pressure and temperature sensors
  - Set of temperature sensors and pressure sensors must be optimally located in a plant
  - Aim to maximize accuracy of temperature and pressure measurements
  - What if location for optimal pressure sensor placement corresponds to location of optimal temperature sensor placement?
  - Difference evaluations determine which sensor will be more beneficial for overall system performance!

Outline

- Motivation
- Critical Concepts
- Project Objectives
- Objective 1: Methodology and Results
- Objective 2: Methodology and Results
- Closing Remarks

Project Objectives

1. Develop performance metrics and algorithms for heterogeneous sensor networks
   - Quantify sensor network effectiveness
   - Allow tradeoffs in communication, computation, and sensing requirements
   - Develop objective functions for sensors (agents)

2. Demonstrate scalability, reconfigurability, and robustness of heterogeneous sensor network
   - Does it work with 10,000 sensors?
   - What if system level goals change?
   - What if sensors fail?

Outline

Defect Combination Problem

- Large set of disparate sensing devices
- Each device has noise and measurement error
- Which subset of devices should be activated for most accurate signal?

$$G = \frac{\left|\sum_{i=1}^{N} n_i a_i\right|}{\sum_{i=1}^{N} n_i}$$

Rankine Cycle Defect Combination Problem

- Apply DCP to each plant state in a Rankine cycle model
- Goal: attain accurate pressure and temperature measurements
- Agent feedback based on work and heat rates

Methodology

- Each agent has a probability distribution regarding which action it selects
- Probability distributions updated via cooperative coevolutionary algorithm
- As evolutionary time progresses, quality of solutions improves

Results: 100 Sensors

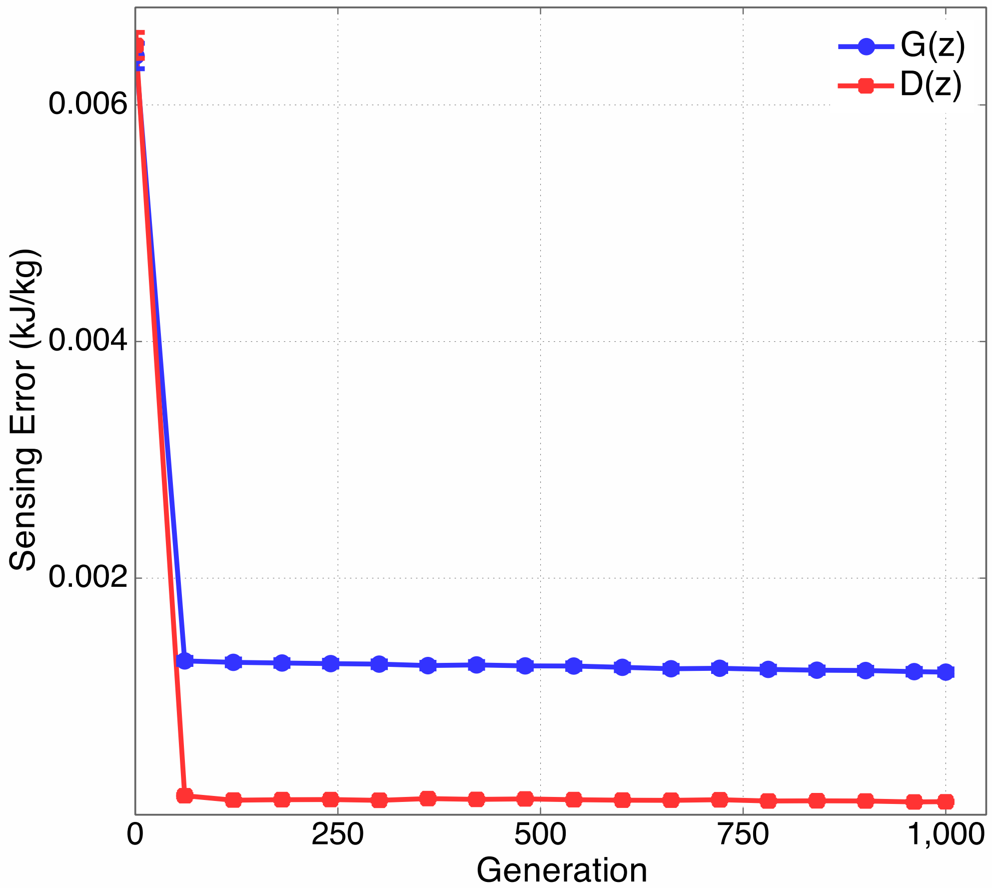

- Difference evaluations result in 9.1% of the error from G(z)

Results: 1000 Sensors

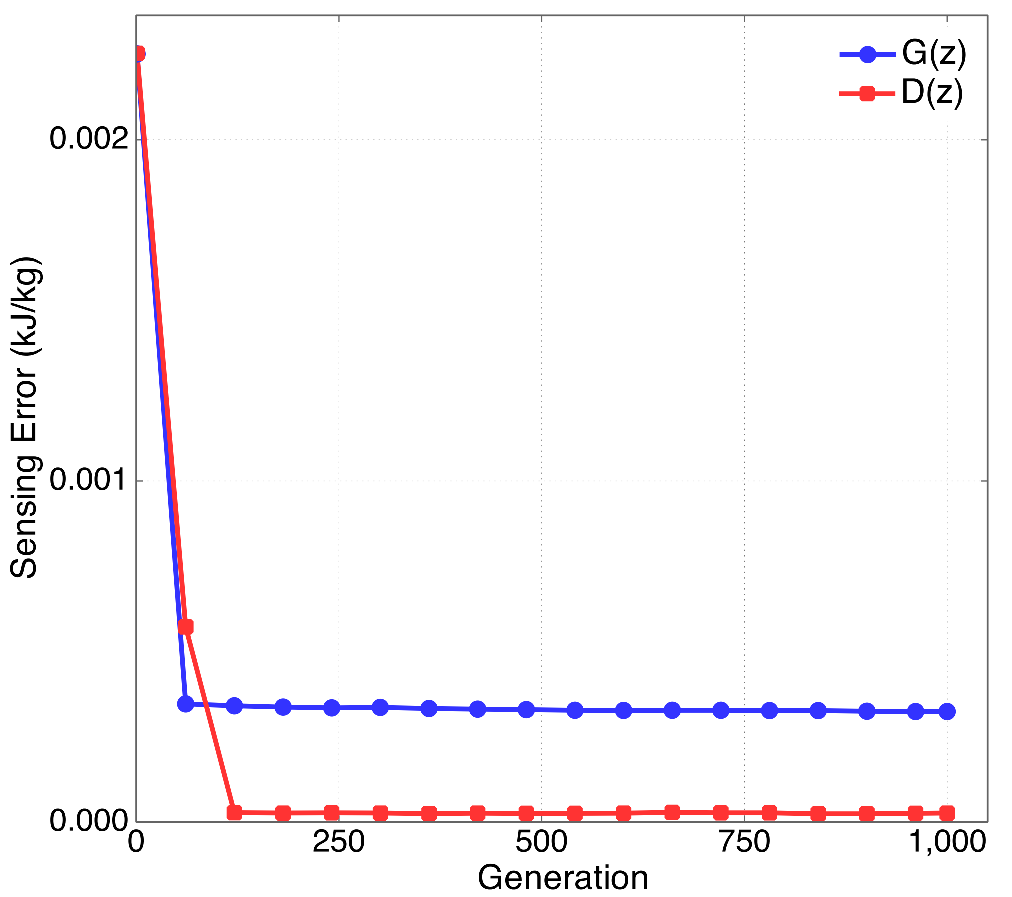

- Difference evaluations result in 1.2% of the error from G(z)

Outline

- Motivation
- Critical Concepts
- Project Objectives
- Objective 1: Methodology and Results
- Objective 2: Methodology and Results
- Closing Remarks

Objective 2

- Scalable: system must scale to thousands of devices
- Reconfigurable: system must adapt to failing devices

Methodology

- Each sensor in the network controlled by a single autonomous agent
- Each agent maintains a Q-table estimating value of sensing
- For each learning step:
  - Agents all take an action
  - Overall system performance computed
  - Agents update Q-tables
- As more learning steps occur, system performance improves

Scalability

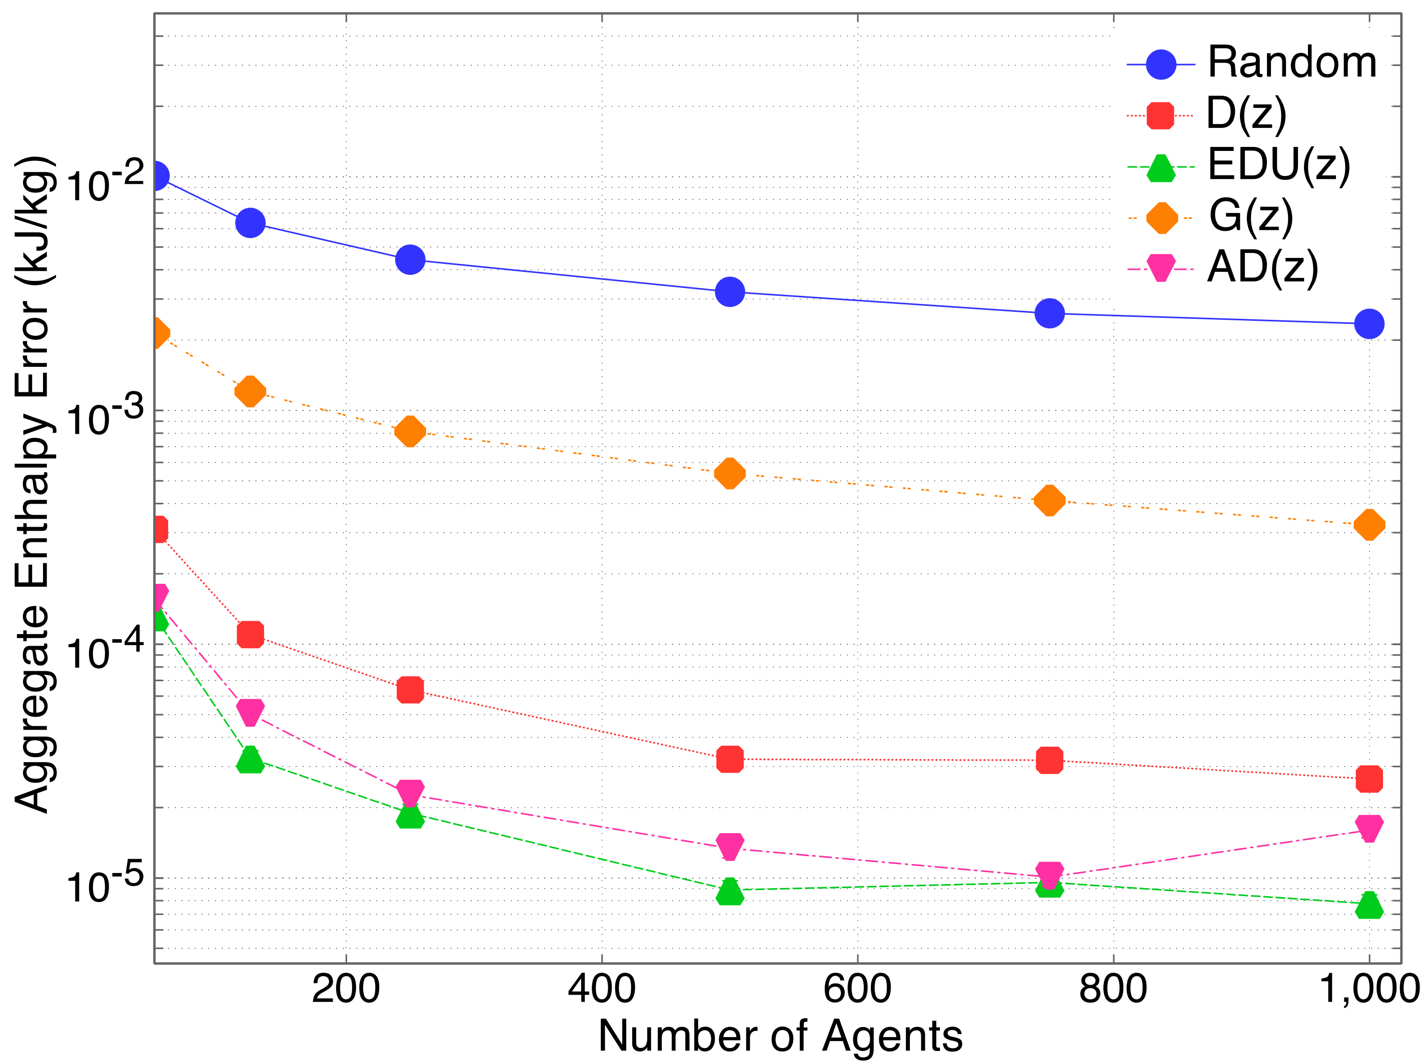

Reconfigurability: 20% Noise, 20% Failures

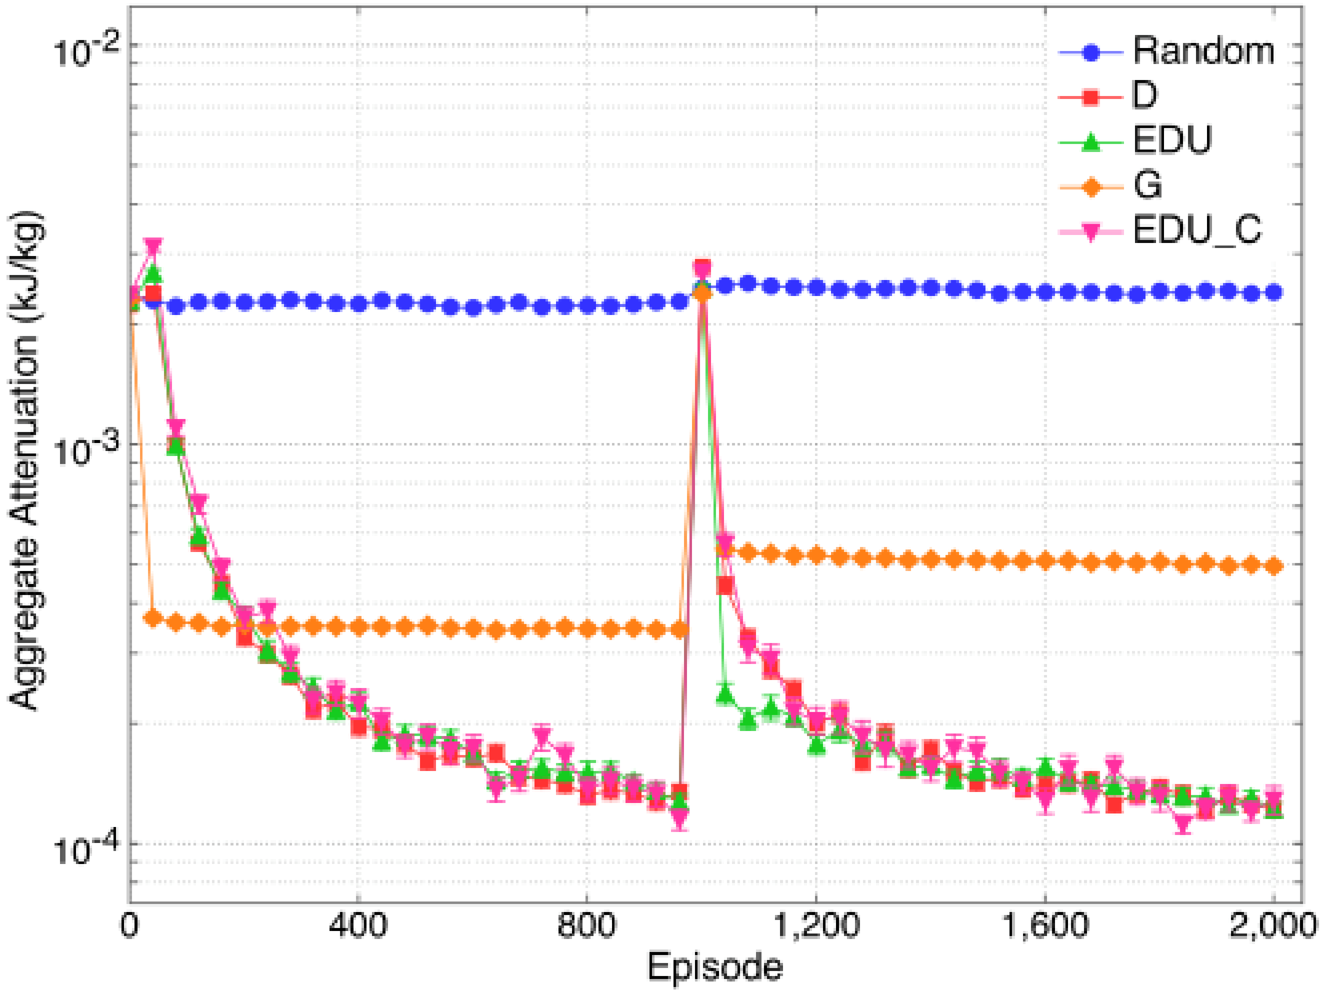

RCDCP: 1,000 Agents, 20% Noise, 20% Failures

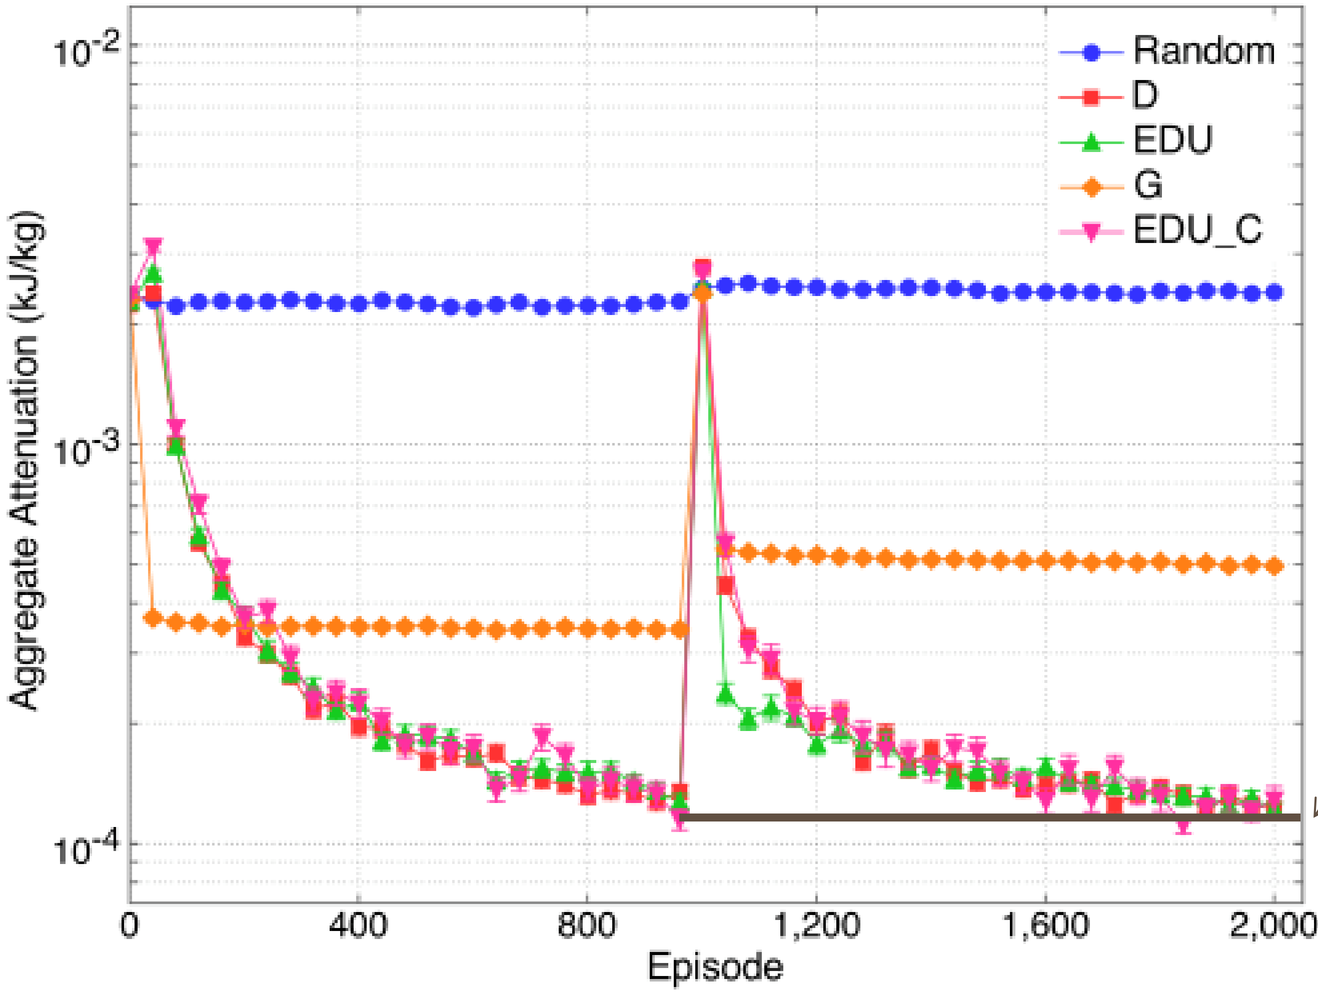

20% failure yields no significant performance drop

Insights

- System is extremely scalable
- System reconfigures with no performance loss after 20% sensor failure
- Network provides extremely accurate measurements, and quickly reconfigures after large changes in system conditions

Outline

- Motivation
- Critical Concepts
- Project Objectives
- Objective 1: Methodology and Results
- Objective 2: Methodology and Results
- Closing Remarks

Closing Remarks

- Proper objective functions improve system performance
- Networks can reconfigure after large disruptions
- Networks are robust to noise
- Networks are extremely scalable

Benefits of Our Approach

- Advanced Energy Systems
  - More efficient information collection
  - Quick response to sudden developments
  - Autonomous system reconfiguration
- Department of Energy and US Government
  - Smart grid
  - Coordinated search and rescue
  - Self-organizing nano/micro devices

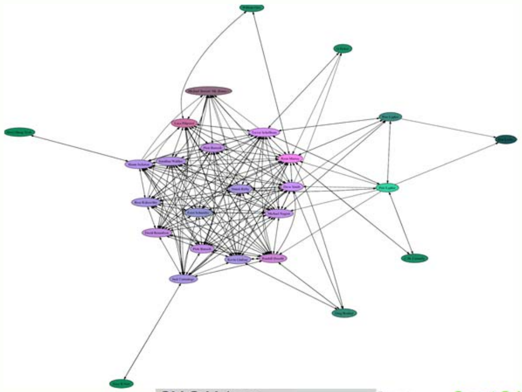

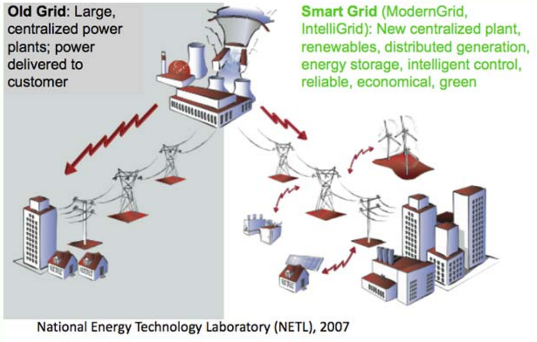

National Energy Technology Laboratory (NETL), 2007

Benefits of Our Approach

- Advanced Energy Systems
  - More efficient information collection
  - Quick response to sudden developments
  - Autonomous system reconfiguration
- Department of Energy and US Government
  - Smart grid
  - Coordinated search and rescue
  - Self-organizing nano/micro devices
- American Public
  - Smart homes
  - Smart highways
  - Smart airports

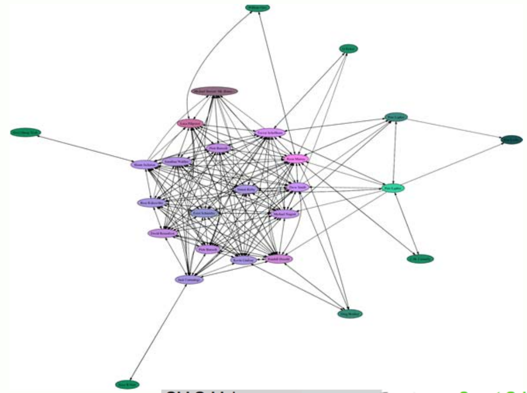

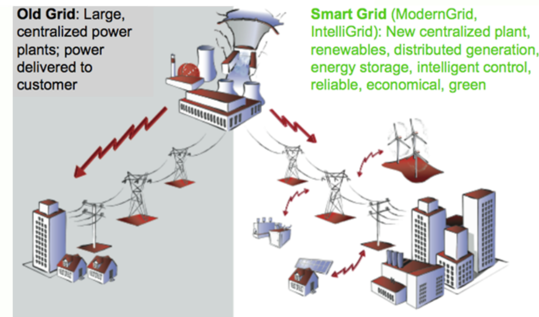

National Energy Technology Laboratory (NETL), 2007

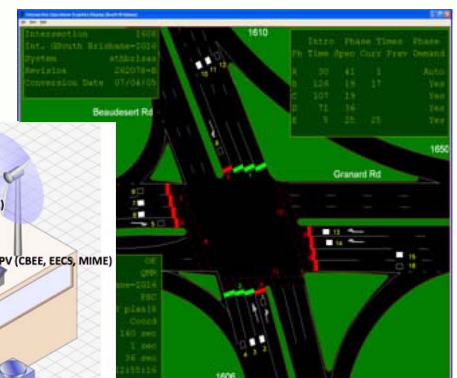

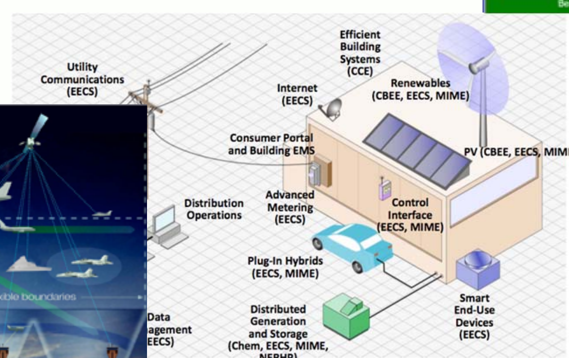

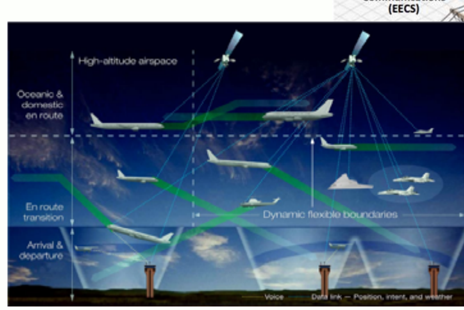

Publications Related to this Research

1. C. Holmes Parker, A. Agogino, and K. Tumer. Evolving distributed resource sharing for cubesat constellations. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Philadelphia, PA, July 2012.
2. C. Holmes Parker, A. Agogino, and K. Tumer. Evolving large scale uav communi- cation systems. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Philadelphia, PA, July 2012. **Best ”Real World Applications” paper award.**
3. M. Colby, C. Holmes Parker, and K. Tumer. Coordination and control for large dis- tributed sensor networks. In *Future of Instrumentation International Workshop (FIIW- 2012)*. Gatlinburg, TN, October 2012.
4. M. Colby and K. Tumer. Multiagent reinforcement learning in a distributed sensor network with indirect feedback. In *In Proceedings of the Twelfth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013)*, Saint Paul, Minnesota.
5. M. Colby and K. Tumer. Performance and fiscal analysis of distributed sensor net- works in a power plant. In *AAMAS-2012 Workshop on Agent Technologies for Energy Systems*. Valencia, Spain, June 2012.
6. C. Holmes Parker and K. Tumer. Combining difference rewards and hierarchies for scaling to large multiagent system. In *AAMAS-2012 Workshop on Adaptive and Learning Agents*. Valencia, Spain, June 2012.
7. C. Roth. Agent objectives for evolving coordinated sensor networks. Master’s thesis, University of Applied Sciences Offenburg, Germany, 2010.
8. C. Roth, M. Knudson, and K. Tumer. Agent fitness functions for evolving coordinated sensor networks. In *Proceedings of the Genetic and Evolutionary Computation Conference*, Dublin, Ireland, July 2011.

Questions?

?

Contact Info:

Kagan Tumer

Oregon State University

kagan.tumer@oregonstate.edu

http://engr.oregonstate.edu/~ktumer