Intelligent Coordination of Heterogeneous Sensors in Advanced Power Systems

PI: Kagan Tumer

Presenter: Mitchell Colby

Oregon State University

kagan.tumer@oregonstate.edu

colbym@engr.orst.edu

Agreement Number: DE-FE0011403

NETL Project Manager: Maria Reidpath

Motivation

- Where are we going?
  - Distributed, complex, hybrid systems
  - Components with higher computational power
- What do we need to account for?
  - 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?

Project Objectives

1. Develop performance metrics and algorithms for heterogeneous sensor network.

2. Demonstrate scalability, reconfigurability, and robustness of heterogeneous sensor network in simulation.

Project Milestones

M1. new objective functions and evolutionary algorithms.

M2. new objective functions and reinforcement learning.

M3. scalability.

M4. reconfigurability and scalability.

Project Milestones

M1. new objective functions and evolutionary algorithms.

M2. new objective functions and reinforcement learning.

Objective 1

M3. scalability.

M4. reconfigurability and scalability.

Objective 2

Project Milestones

M1. new objective functions and evolutionary algorithms.

M2. new objective functions and reinforcement learning.

M3. scalability.

M4. reconfigurability and scalability.

Milestone 1

- Evolutionary algorithms
- New objective function

Evolutionary Algorithms (Single Sensor)

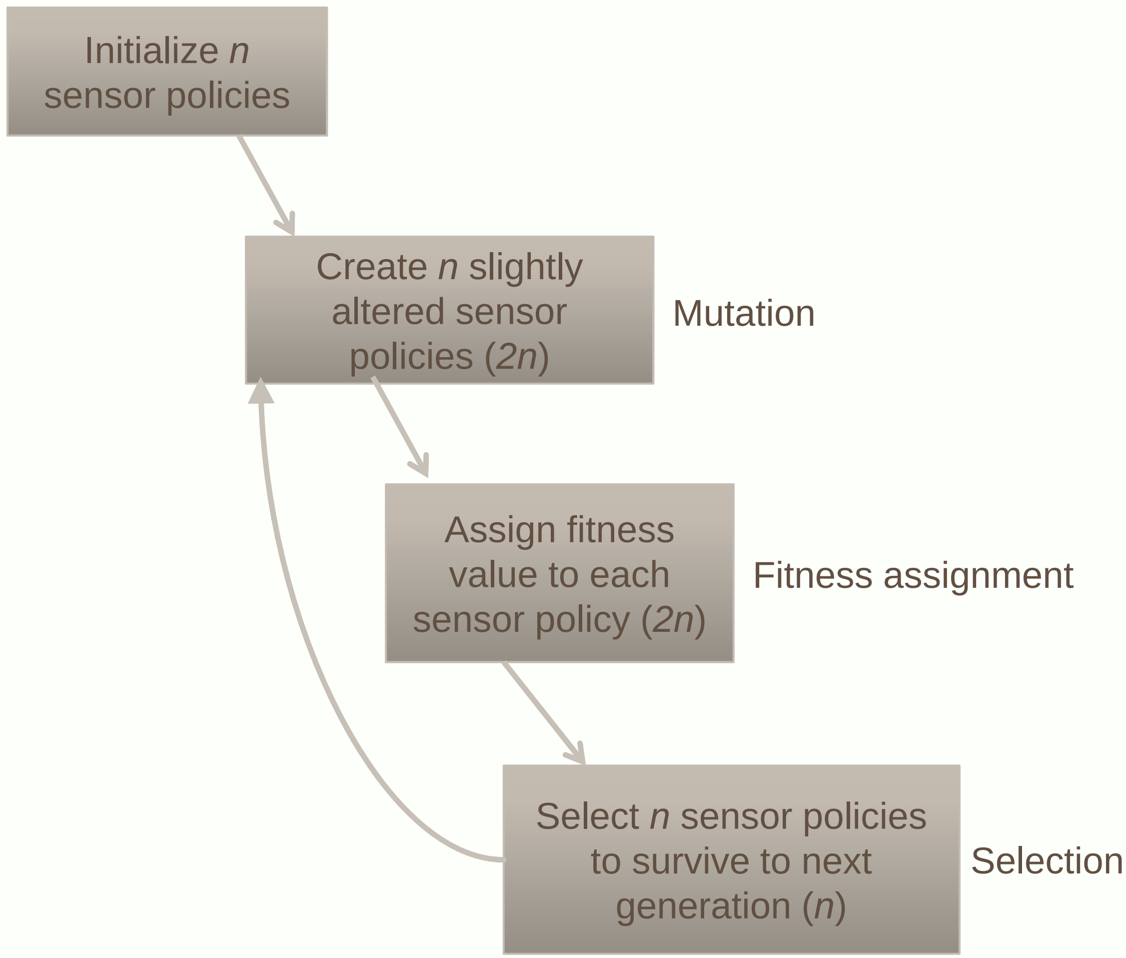

What About Multiple Sensors?

- Extend evolutionary algorithms for multiagent systems.
- Cooperative coevolution: multiple parallel EAs.
- Fitness assignment is based on:
  - Agent's policy
  - Collaborating agents' policies

Global Evaluation Functions

- Assign fitness using team performance
- Too much noise (not ***sensitive***)

- Example: 100 agents
  - 99 agents perform optimally
  - 1 agent does nothing
  - Fairly high system performance $G(z)$

Local Evaluation Functions

- Assign fitness based on local measures
- Greedy agents (not ***aligned***)

- Example: Tragedy of the Commons
  - Agents overuse shared resources, and hurt system

Desirable Fitness Function Properties

	Aligned	Sensitive
Global	✔	
Local		✔

Desirable Fitness Function Properties

	Aligned	Sensitive
Global	✔	
Local		✔
???	✔	✔

Difference Evaluation Functions

- Difference Evaluation Functions:

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

Difference Evaluation Functions

- Difference Evaluation Functions:

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

- Second term: removes noise caused by other agents.

Difference Evaluation Functions

- Difference Evaluation Functions:

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

- Second term: removes noise caused by other agents.

Sensitive ✔

Difference Evaluation Functions

- Difference Evaluation Functions:

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

- Second term: removes noise caused by other agents.

Sensitive ✔

- Derivatives:

$$\frac{\partial D_i(z)}{\partial a_i} = \frac{\partial G(z)}{\partial a_i} - \frac{\partial G(z_{-i} + c_i)}{\partial a_i}$$

Difference Evaluation Functions

- Difference Evaluation Functions:

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

- Second term: removes noise caused by other agents.

Sensitive ✔

- Derivatives:

$$\frac{\partial D_i(z)}{\partial a_i} = \frac{\partial G(z)}{\partial a_i} - \cancel{\frac{\partial G(z_{-i} + c_i)}{\partial a_i}}$$

Difference Evaluation Functions

- Difference Evaluation Functions:

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

- Second term: removes noise caused by other agents.

Sensitive ✔

- Derivatives:

$$\frac{\partial D_i(z)}{\partial a_i} = \frac{\partial G(z)}{\partial a_i} - \cancel{\frac{\partial G(z_{-i} + c_i)}{\partial a_i}}$$

$$\frac{\partial D_i(z)}{\partial a_i} = \frac{\partial G(z)}{\partial a_i}$$

Difference Evaluation Functions

- Difference Evaluation Functions:

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

- Second term: removes noise caused by other agents.

Sensitive ✔

- Derivatives:

$$\frac{\partial D_i(z)}{\partial a_i} = \frac{\partial G(z)}{\partial a_i} - \cancel{\frac{\partial G(z_{-i} + c_i)}{\partial a_i}}$$

$$\frac{\partial D_i(z)}{\partial a_i} = \frac{\partial G(z)}{\partial a_i}$$

Aligned ✔

Approach

- Optimize sensor network performance with CCEAs
- Assign fitness using difference evaluations

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?

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

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

Results: 100 Sensors

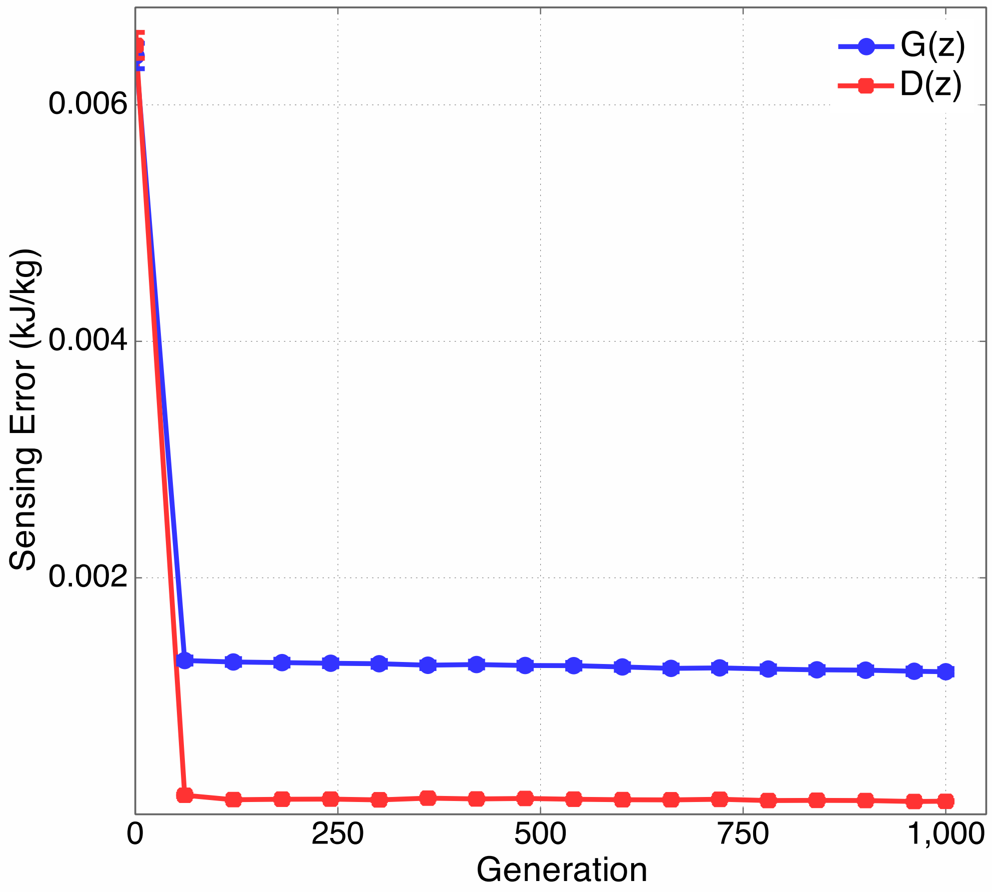

- Difference evaluations result in 9.1% of the error from G(z)

Results: 1000 Sensors

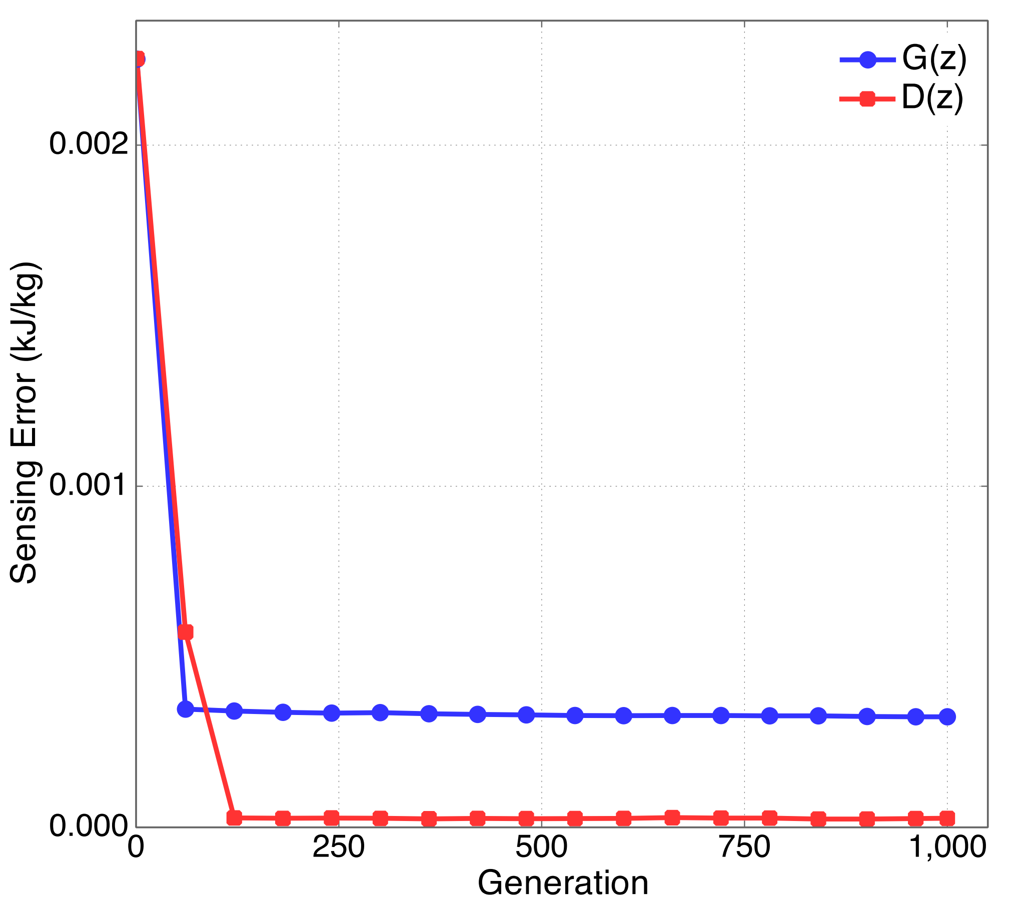

- Difference evaluations result in 1.2% of the error from G(z)

Project Milestones

M1. new objective functions and evolutionary algorithms.

M2. new objective functions and reinforcement learning.

M3. scalability.

M4. reconfigurability and scalability.

Milestone 2

- Multiagent reinforcement learning
- New objective functions

Multiagent Reinforcement Learning

- Multiagent Reinforcement Learning:
  - Maintain expected value for each action
  - Update expected value after taking action and receiving reward

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

Multiagent Reinforcement Learning

- Multiagent Reinforcement Learning:
  - Maintain expected value for each action
  - Update expected value after taking action and receiving reward

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

- Intuition: actions with high rewards are reinforced
  - Think Pavlov

Multiagent Reinforcement Learning

- Multiagent Reinforcement Learning:
  - Maintain expected value for each action
  - Update expected value after taking action and receiving reward

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

- Intuition: actions with high rewards are reinforced
  - Think Pavlov
- Rewards:
  - based on team's performance
  - same problems with alignment and sensitivity!

New Agent Objective Functions

- Difference Evaluation Function:

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

New Agent Objective Functions

- Difference Evaluation Function:

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

- Expected Difference Evaluation Function:

$$ED_i(z) = G(z) - \sum_j P(c_j) G(z_{-i} + c_j)$$

New Agent Objective Functions

- Difference Evaluation Function:

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

- Expected Difference Evaluation Function:

$$ED_i(z) = G(z) - \sum_j P(c_j) G(z_{-i} + c_j)$$

- Average Difference Evaluation Function:

$$AD_i(z) = G(z) - \frac{1}{N} \sum_{j=1}^{N} G(z_{-i} + c_j)$$

Oregon State UNIVERSITY

Approach

- Optimize sensor network performance with reinforcement learning
- Assign rewards using difference evaluation variants
- Rankine cycle DCP

Results: 100 Agents

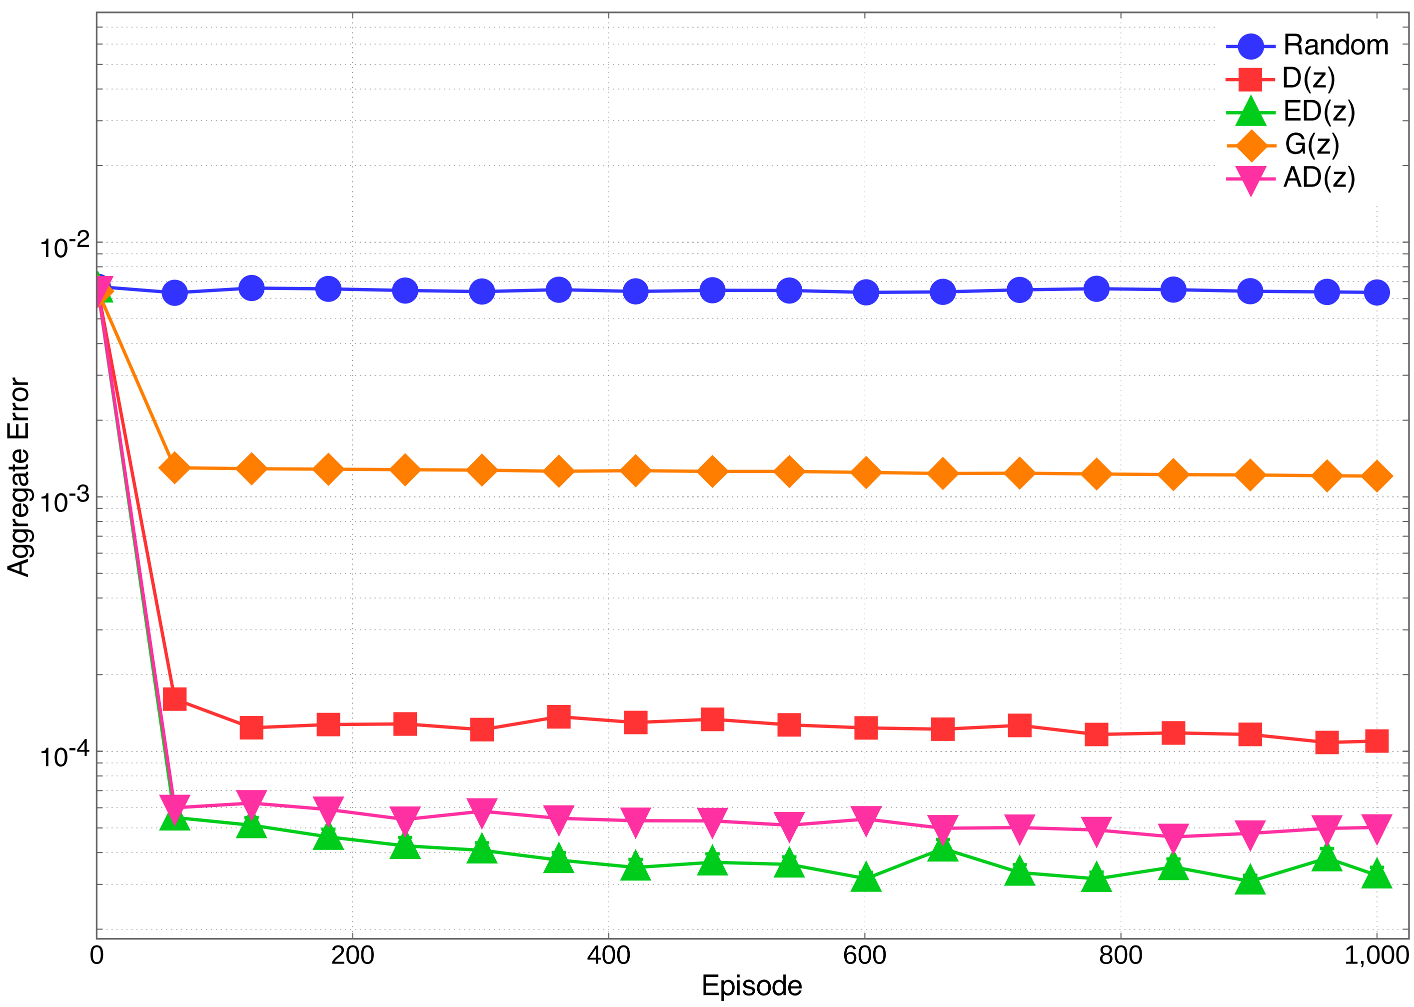

Results: 1000 Agents

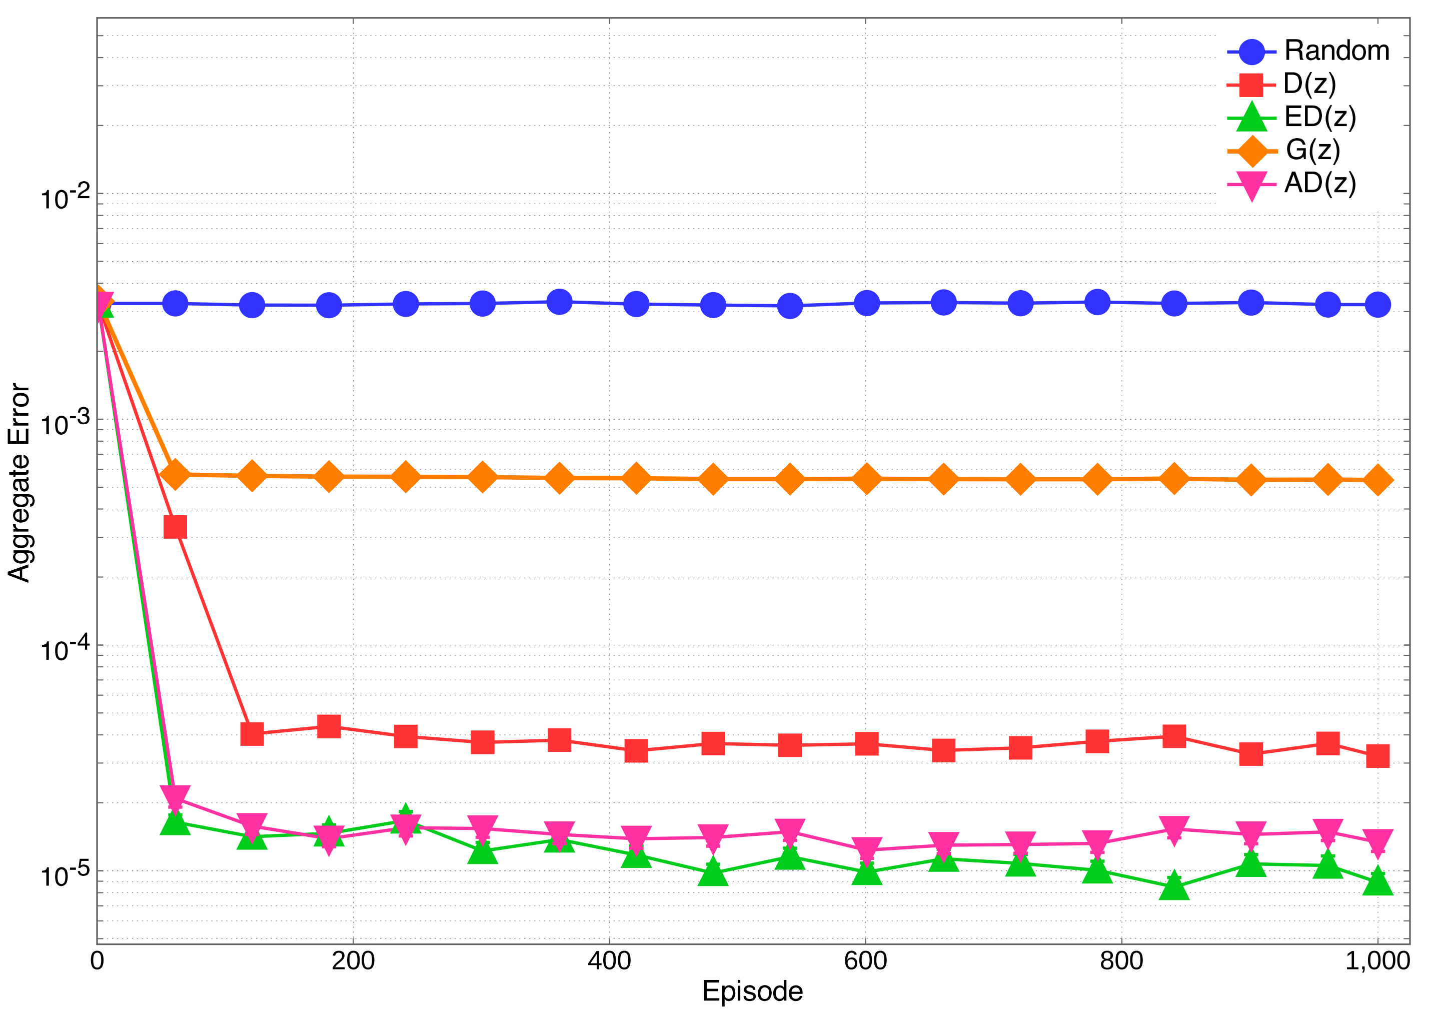

Project Milestones

M1. new objective functions and evolutionary algorithms.

M2. new objective functions and reinforcement learning.

M3. scalability.

M4. reconfigurability and scalability.

Milestone 3

- Demonstrate system scalability
- What about 2000 sensors?

Approach

- Optimize sensor network performance with reinforcement learning
- Assign fitness using difference evaluation variants
- Rankine cycle DCP

Results: 2000 Agents

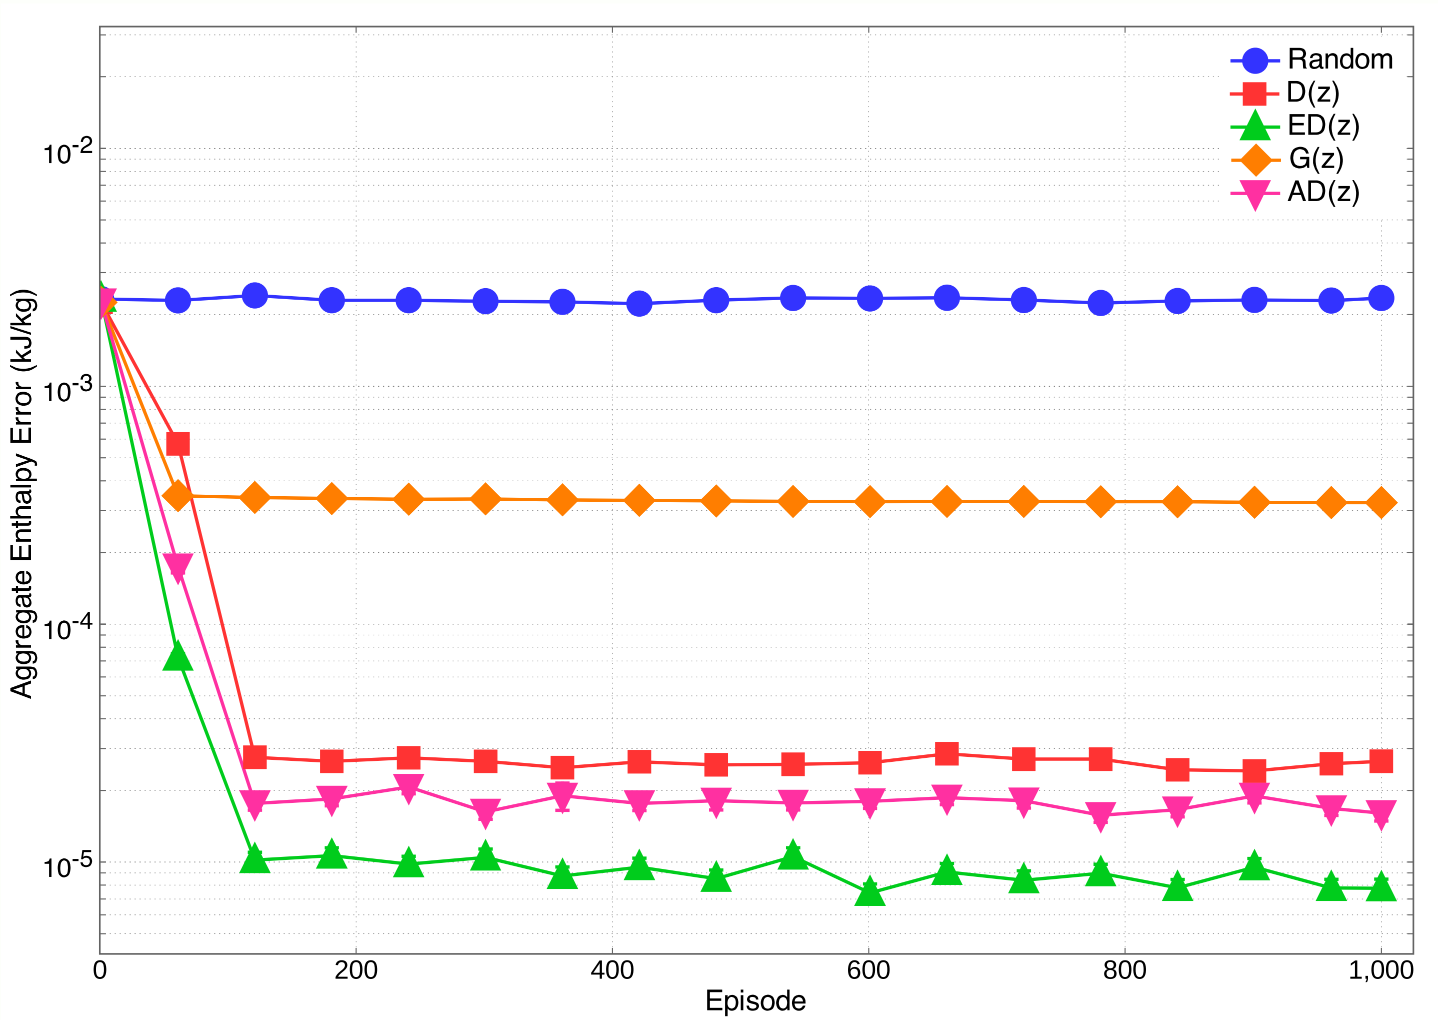

Results: Scalability

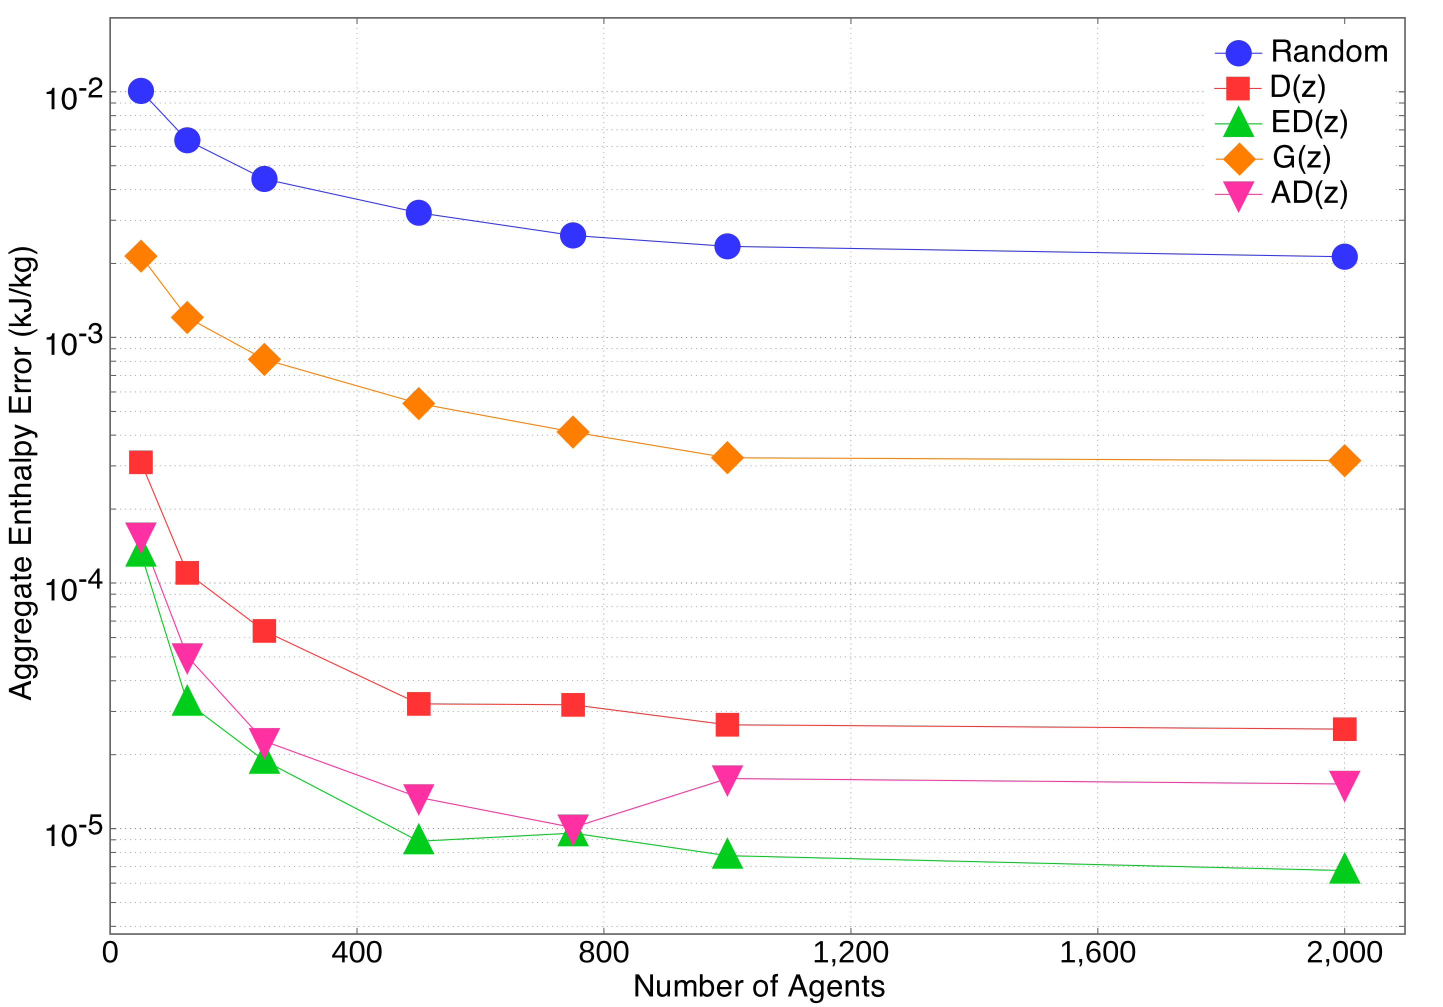

Mitchell Colby, Kagan Tumer (PI), Oregon State University

Project Milestones

M1. new objective functions and evolutionary algorithms.

M2. new objective functions and reinforcement learning.

M3. scalability.

M4. reconfigurability and scalability.

Milestone 4

- Reconfigurability
  - robustness to noise
  - system reconfiguration after device failure
- Scalability

Experiments

- Add noise to system
- Agent (sensor) failures

Experiments

- Add noise to system
- Agent (sensor) failures

1000 Agents, 10% Sensor Noise

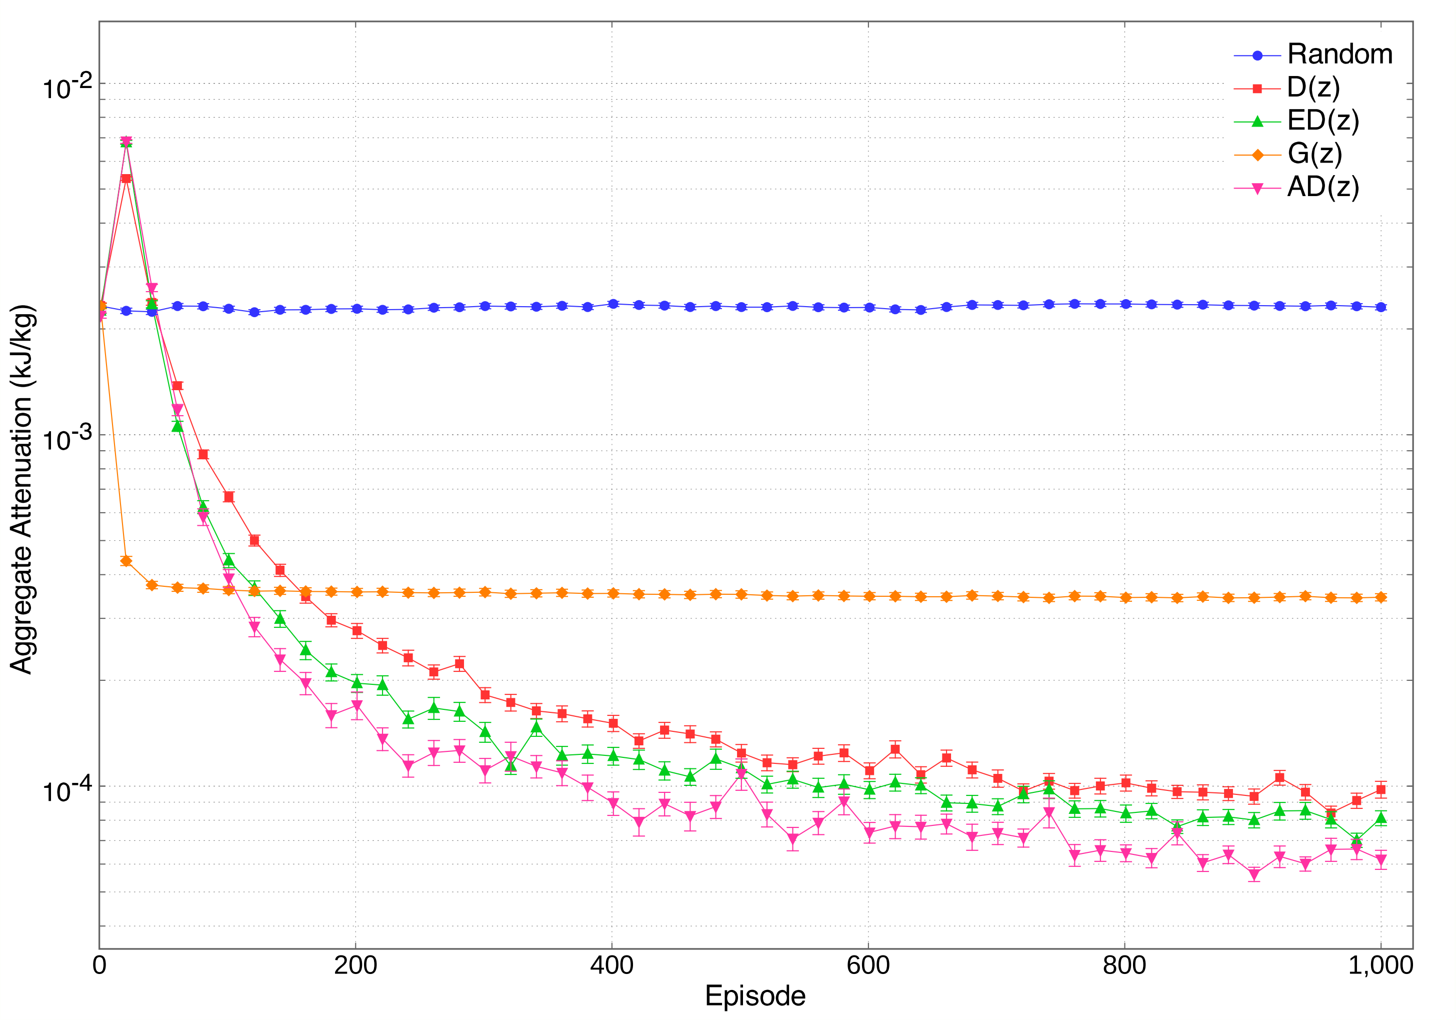

Mitchell Colby, Kagan Tumer (PI), Oregon State University

1000 Agents, 50% Sensor Noise

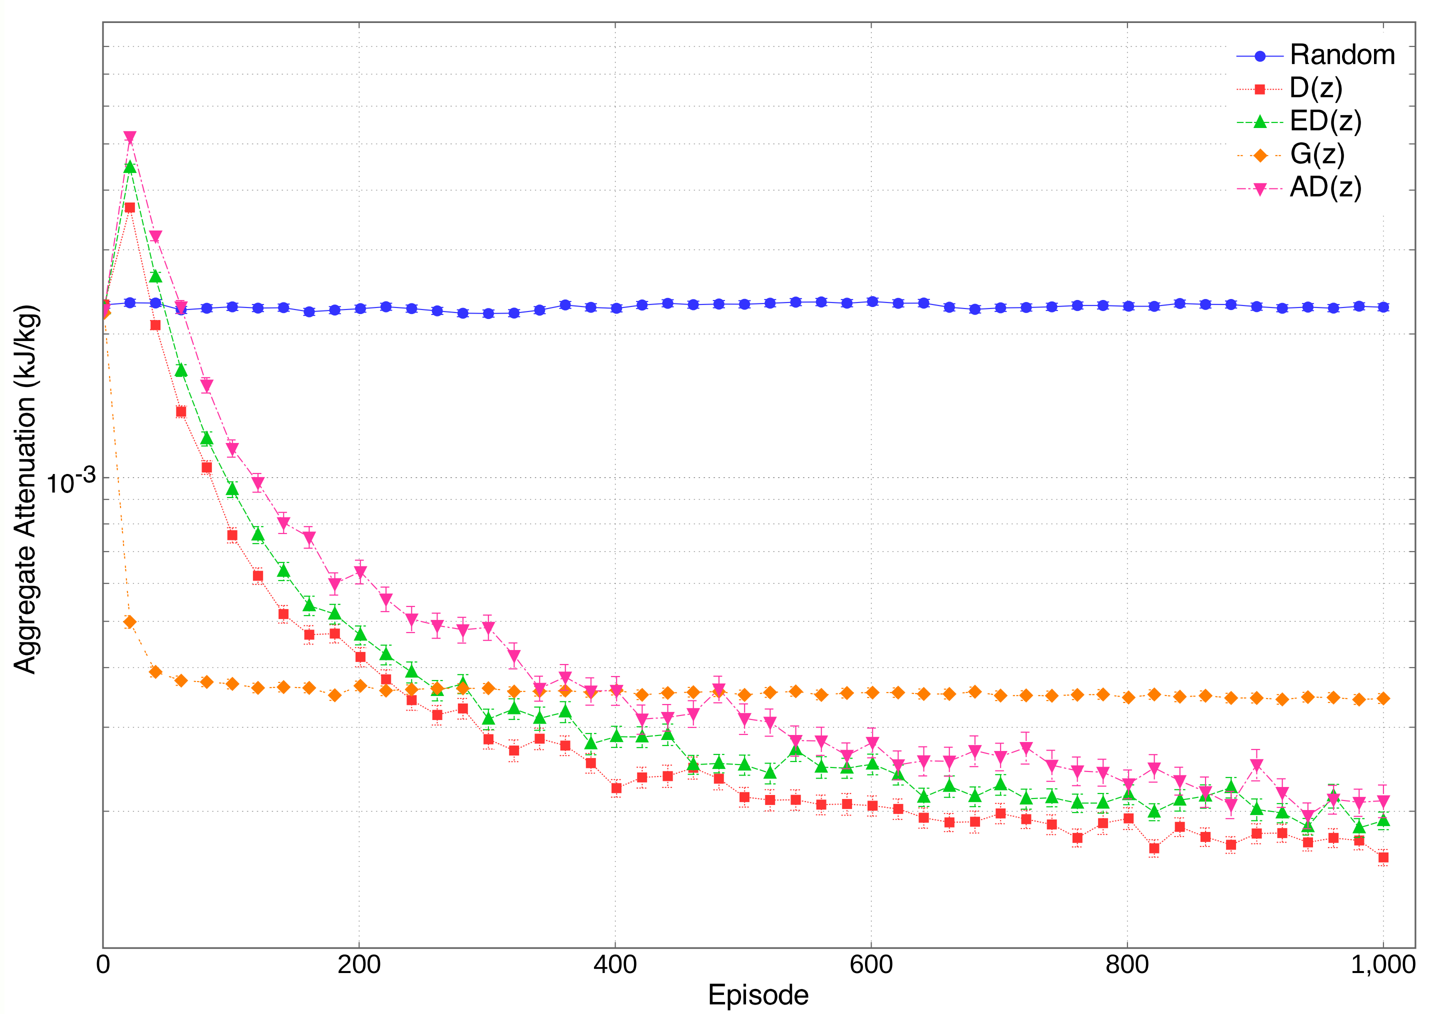

Experiments

- Add noise to system
- Agent (sensor) failures

1000 Agents, 10% Noise, 10% Failures

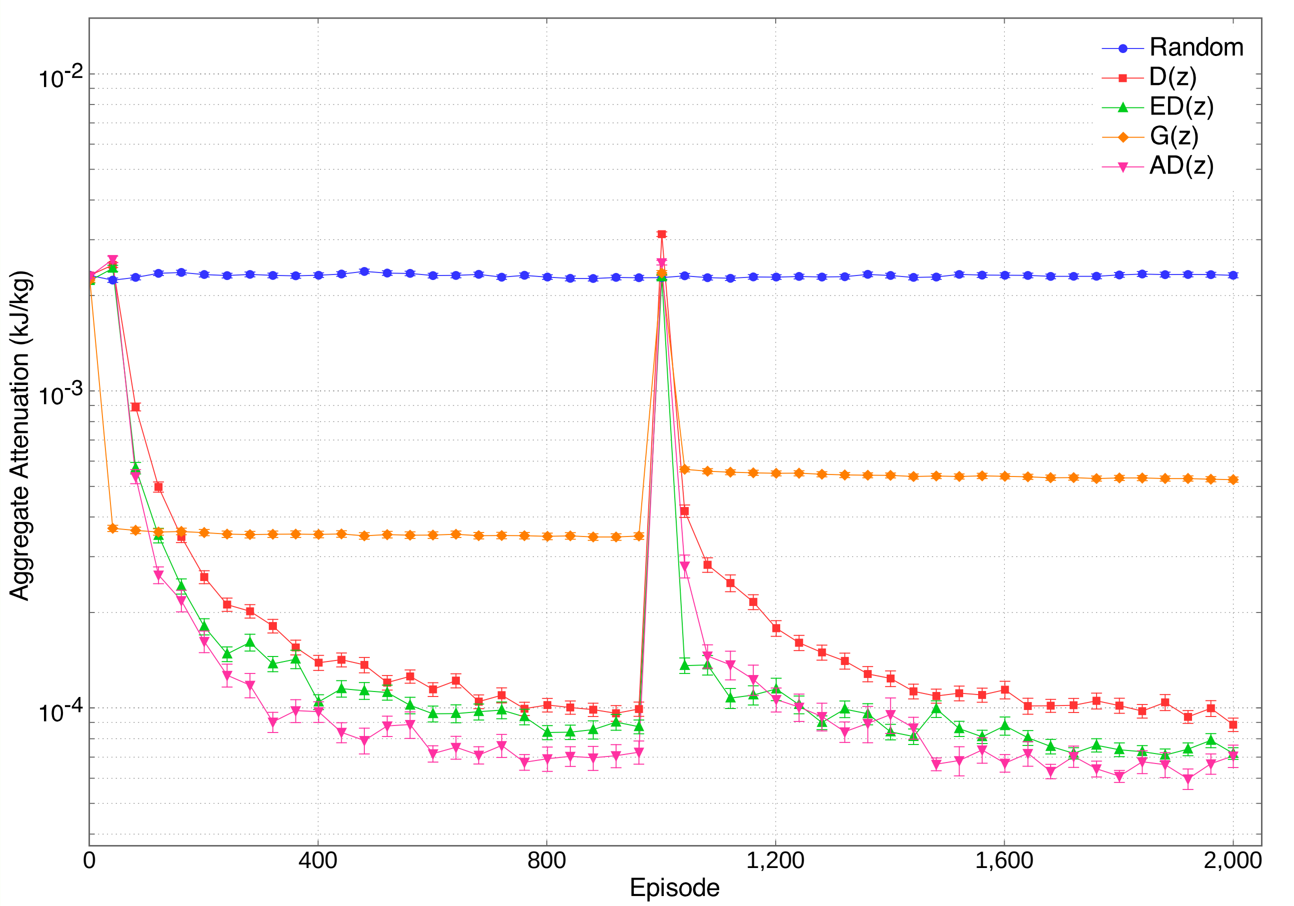

Mitchell Colby, Kagan Tumer (PI), Oregon State University

1000 Agents, 10% Noise, 10% Failures

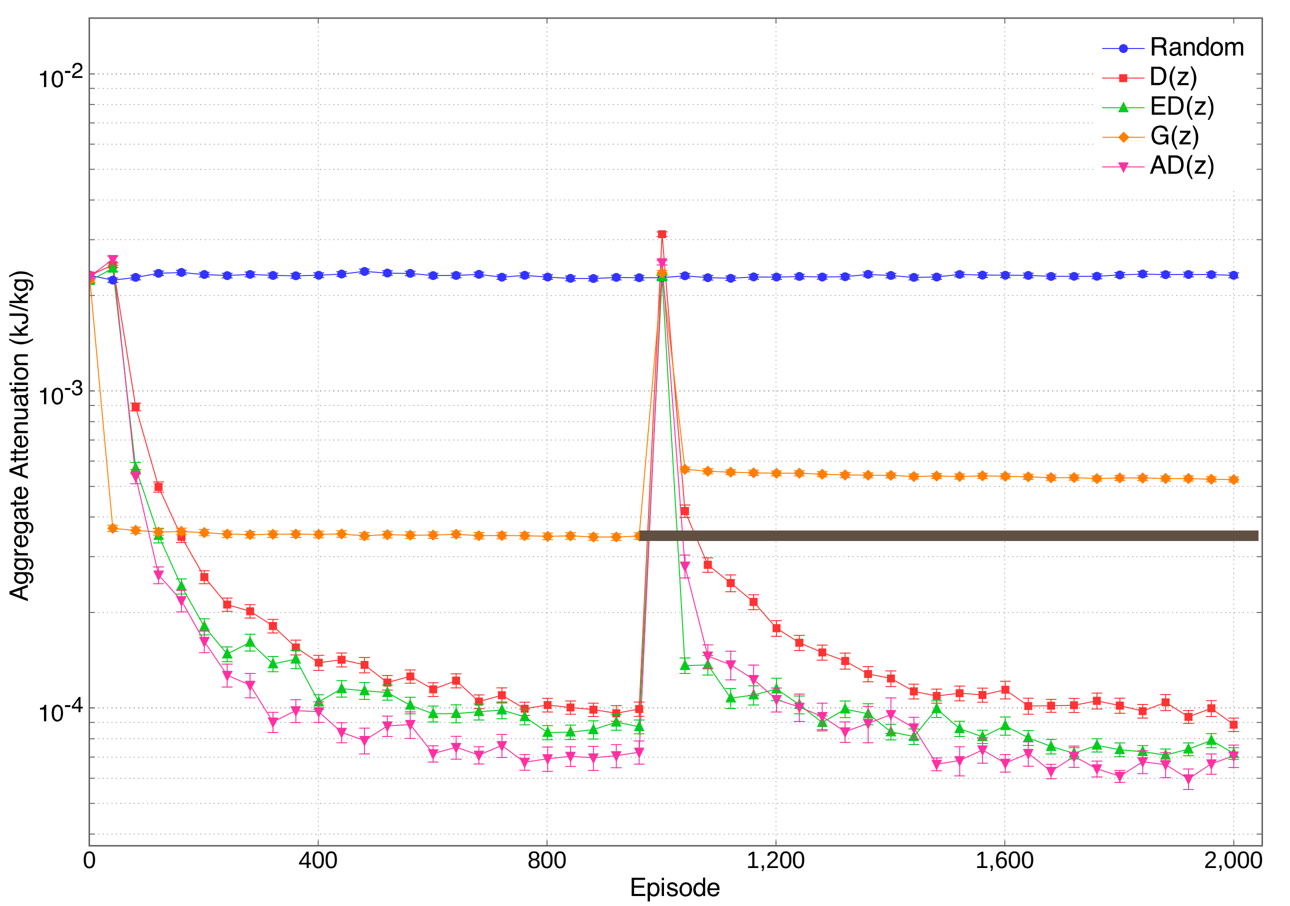

10% loss of sensors leads to loss of system performance

1000 Agents, 10% Noise, 10% Failures

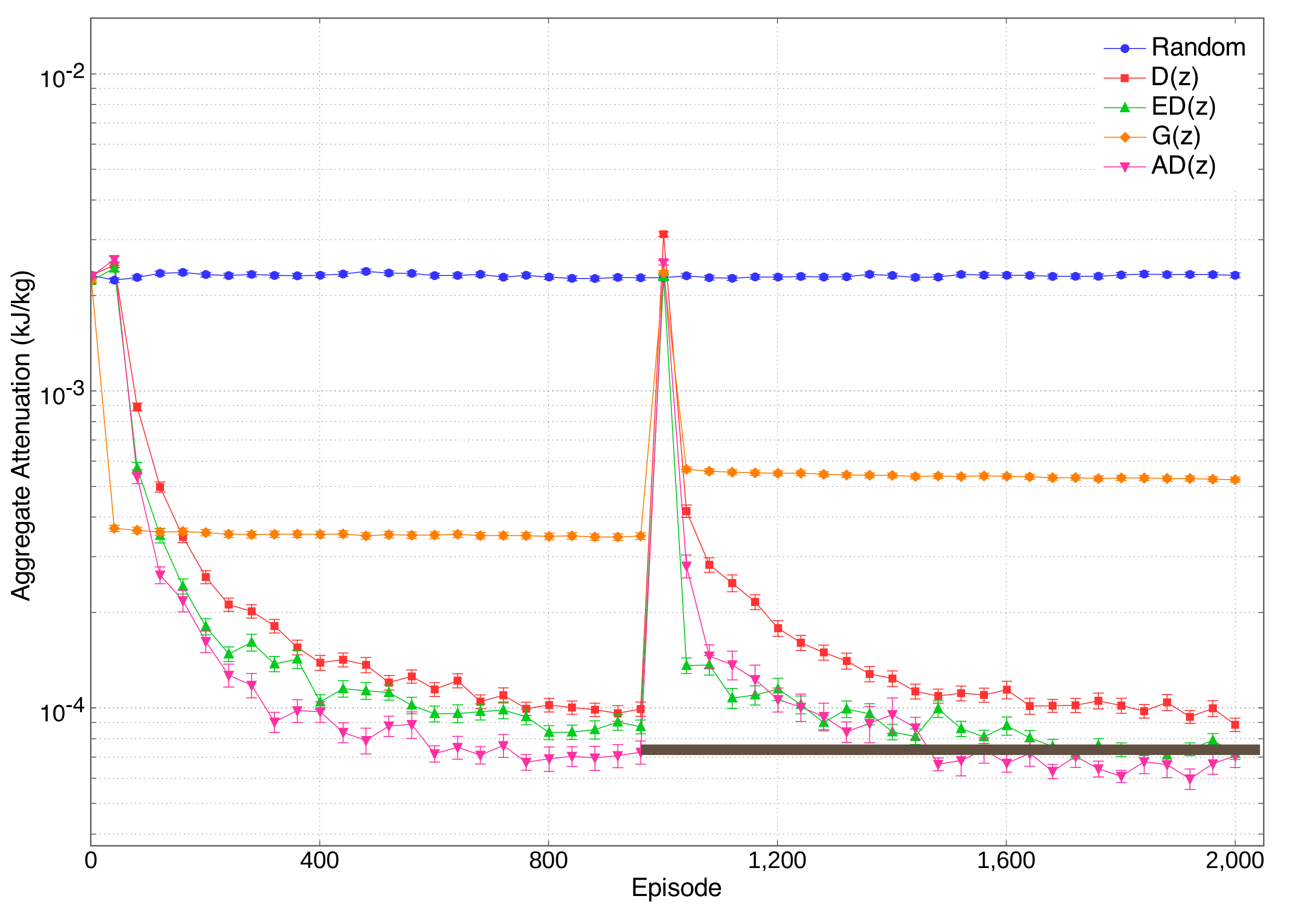

10% loss of sensors leads to no loss of system performance

Mitchell Colby, Kagan Tumer (PI), Oregon State University

2000 Agents, 50% Noise, 50% Failure

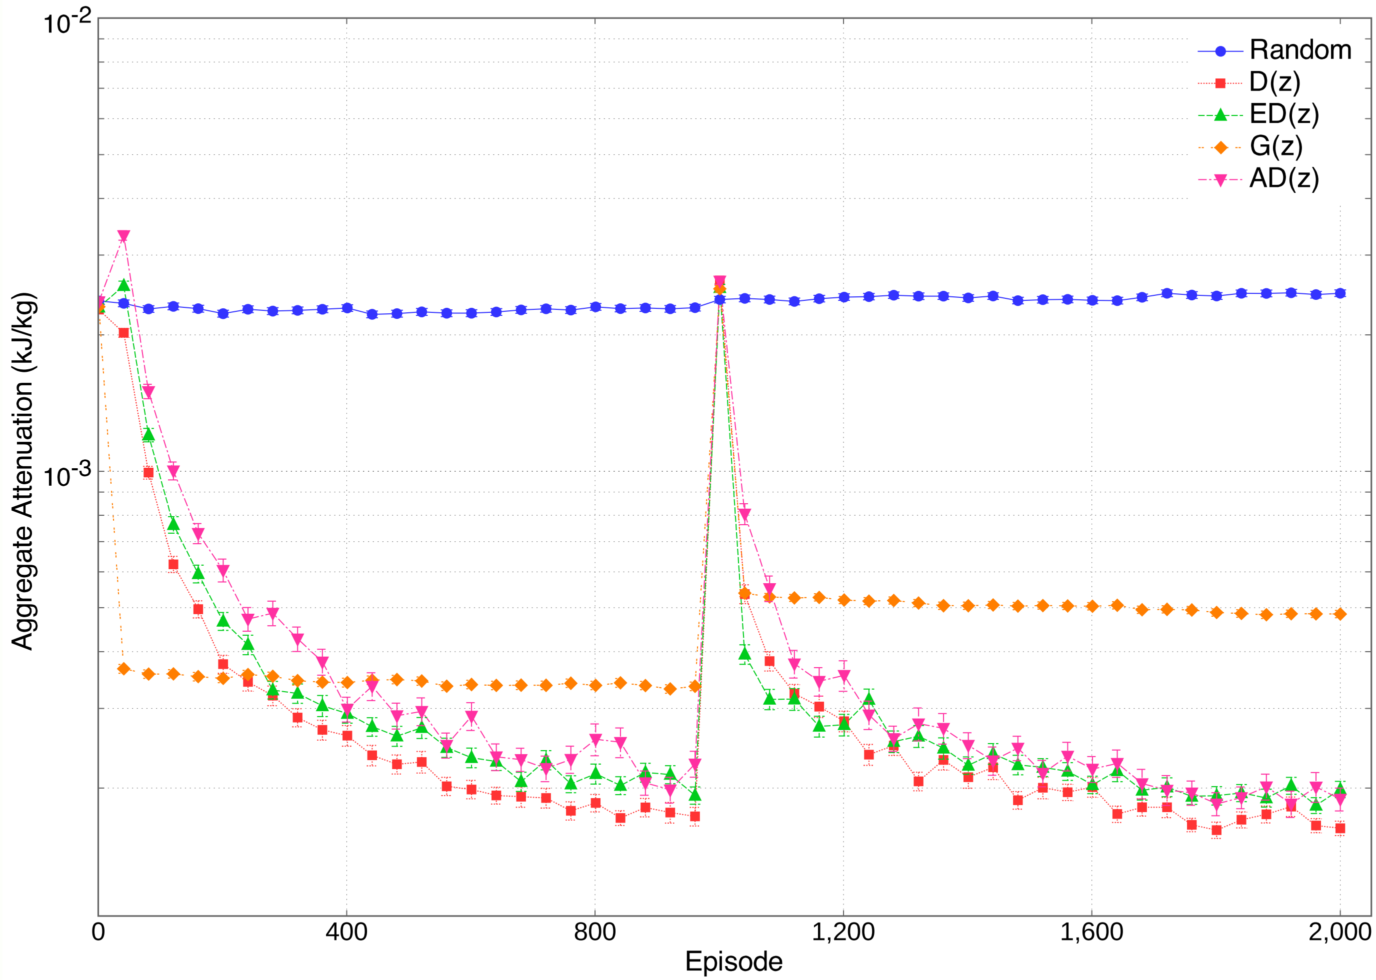

Closing Remarks

- Difference objective functions improve system performance
- EAs vs RL: timescale

- Sensor networks:
  - can reconfigure after large disruptions
  - are robust to noise
  - are extremely scalable

Publications Related to this Research

1. M. Colby, S. Kharaghani, C. HolmesParker, and K. Tumer. Counterfactual Exploration for Improving Multiagent Learning. *In Proceedings of the Fourteenth International Joint Conference On Autonomous Agents and Multiagent Systems (AAMAS 2015), Istanbul, Turkey.*
2. M. Colby and K. Tumer. Learning-Based Coordination of Large Distributed Sensor Networks. In *Proceedings of the 2015 International Society of Automation Power Industry Division Symposium,* Kansas City, MO.
3. M. Colby, W. Curran, C. Rebhuhn, and K. Tumer. Approximating Difference Evaluations with Local Knowledge. *In Proceedings of the Thirteenth International Joint Conference On Autonomous Agents and Multiagent Systems (AAMAS 2015), Paris, France.*
4. M. Colby and K. Tumer. Distributed Sensor Network Control with Evolutionary Algorithms. In *Proceedings of the 2014 International Society of Automation Power Industry Division Symposium,* Scottsdale, AZ.
5. M. Colby. Theoretical and Implementation Improvements for Difference Evaluation Functions. Ph.D. Dissertation, Oregon State University, June 2014.
6. A. Rahmattalabi, M. Colby, and K. Tumer. Evolving Control Policies for Distributed Sensor Network Coordination. In *Proceedings of the 2015 International Society of Automation Power Industry Division Symposium,* Kansas City, MO.

Students Supported by this Grant

- Mitchell Colby (Ph.D., 2014)
- Logan Yliniemi (Ph.D., June 2015)
- Drew Gabler (M.S., June 2015)
- Sepideh Kharaghani (M.S., June 2015)
- Jacob Hamar (Undergraduate Assistant)

Acknowledgements

- DOE NETL
- Project Manager: Maria Reidpath
- Students: Sepideh Kharaghani, Aida Rahmattalabi

Questions?

?

PI: Kagan Tumer, Oregon State University

kagan.tumer@oregonstate.edu

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

Presenter: Mitchell Colby, Oregon State University

colby.mk@gmail.com

http://www.mitchellcolby.com