

Collective Intelligence

PI: David Wolpert

Co-I: Kagan Tumer

NASA Ames Research Center

kagan@ptolemy.arc.nasa.gov

<http://ic.arc.nasa.gov/~kagan>

<http://ic.arc.nasa.gov/projects/COIN/index.html>



What the **** is a Collective?

- A **collective** is a system:
 - With a **world utility** function which measures the full system's performance
 - Composed of many **agents**
 - Where each agent has a **private utility** it is trying to optimize
- Important issues:
 - *How should one set private utility functions?*
 - *How should one update them (team formation)?*
 - *How should utilities be modified in presence of communication restrictions?*



An Analogy: A Company

- World utility ↔ Valuation of company
- Agents ↔ Employees
- Private Utilities ↔ Compensation packages
- Design problem (faced by the board):
 - *How to set/modify compensation packages (private utilities) of the agents to increase valuation of company (world utility)*
 - *Salary/bonus*
 - *Benefits*
 - *Stock options*
 - Note: Board does not tell each individual what to do. They set the “incentive packages” for employees (including the CEO).



Collectives of Interest to IS

- ✓ – Control of a constellation of communication satellites
- ✓ – Routing data/vehicles over a communication network/highway
- Dynamic data migration over large distributed databases
- Dynamic job scheduling across a (very) large computer grid
- ✓ – Coordination of rovers/submersibles on Mars/Europa
- Control of the elements of an amorphous telescope
- ✓ – Construction of distributed algorithms for optimization
- ✓ – Selection of components to minimize aggregate error
- Compilation in randomly assembled nanocomputers

Collective intelligence is an enabling technology



Key Concepts for Collectives

- **Factoredness:** Degree to which an agent's private utility is "aligned" with the world utility
 - e.g. stock options are factored w.r.t. company valuation.
 - **Learnability:** Based on sensitivity of an agent's private utility to changes in its state (signal-to-noise).
 - e.g., performance bonuses increase learnability of agent's utility
-
- Interesting question: If you could, would you want everyone's utility to be valuation of company?
 - Factored, yes; but what about learnability?



Brief Illustration of Theory

- Our ability to control system consists of setting some parameters s (e.g, compensation packages):

$$P(G|s)$$



Explore vs. Exploit

Operations Research
Search



Factoredness

Economics
Mechanism Design

Learnability

Machine Learning
Computer Science

- ϵ_G and ϵ_g are intelligences for the agents w.r.t the world utility (G) and their private utilities (g) , respectively

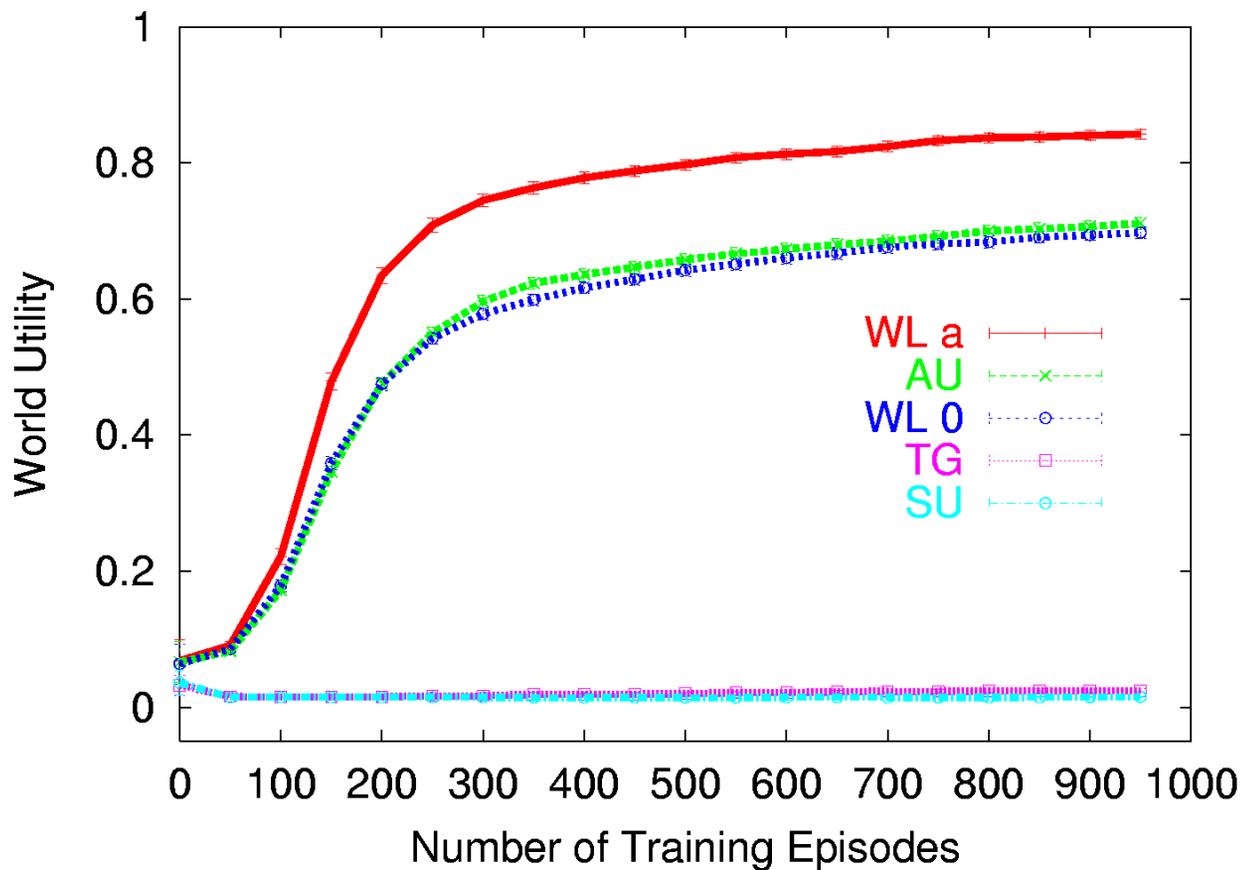


Current/Future Projects

- Application Domains:
 - Multi-rover coordination (Tumer, Agogino) 🕶️
 - Distributed optimization (Wolpert, Tumer) 🕶️
 - Dynamic job scheduling (Tumer, Lawson) 🕶️
 - Distributed resource allocation 1 (Wolpert, Tumer, Aireau) 🕶️
 - Autonomous defect problem (Wolpert, Tumer) 🕶️ 🌶️
 - Nanocomputer compilation (Wolpert, Millonas) 🕶️ 🌶️
 - Distributed resource allocation 2 (Tumer, Agogino) 🌶️
- Scientific Issues:
 - Communication restrictions 🕶️ 🌶️
 - Team formation 🌶️
 - Factoredness/Learnability trade-offs 🕶️ 🌶️



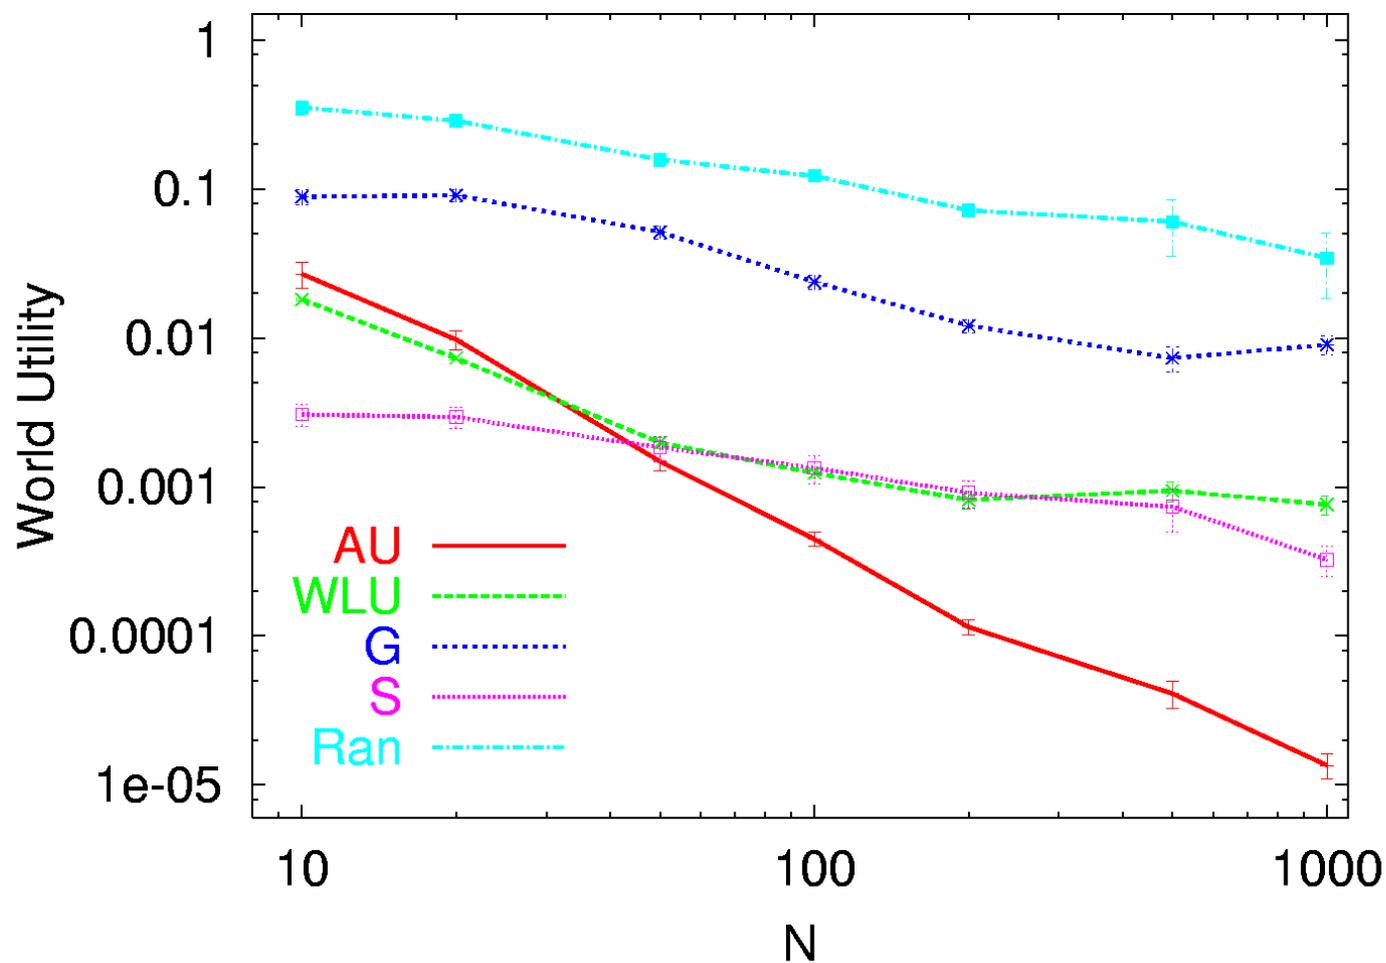
Rover Problem: Utility Comparison



100 rovers on a 32x32 grid



Autonomous Defects Problem: Scaling





Summary



- A collective is a set of “selfish” agents pursuing their own private utilities along with a world utility rating performance of full system.
- Theory of collectives shows *how to configure and/or update the private utilities of the agents so that they “unintentionally cooperate” to optimize the world utility*
- Private utilities based on this theory successfully applied to many domains (e.g., autonomous rovers, constellations of communication satellites, data routing, autonomous defects)
- Associated improvement in performance **increase** with size of problem
- **A fully mature “science of collectives” would benefit the IS project and enable many NASA applications**

THEORY DETAILS:

Nomenclature

Aristocrat Utility

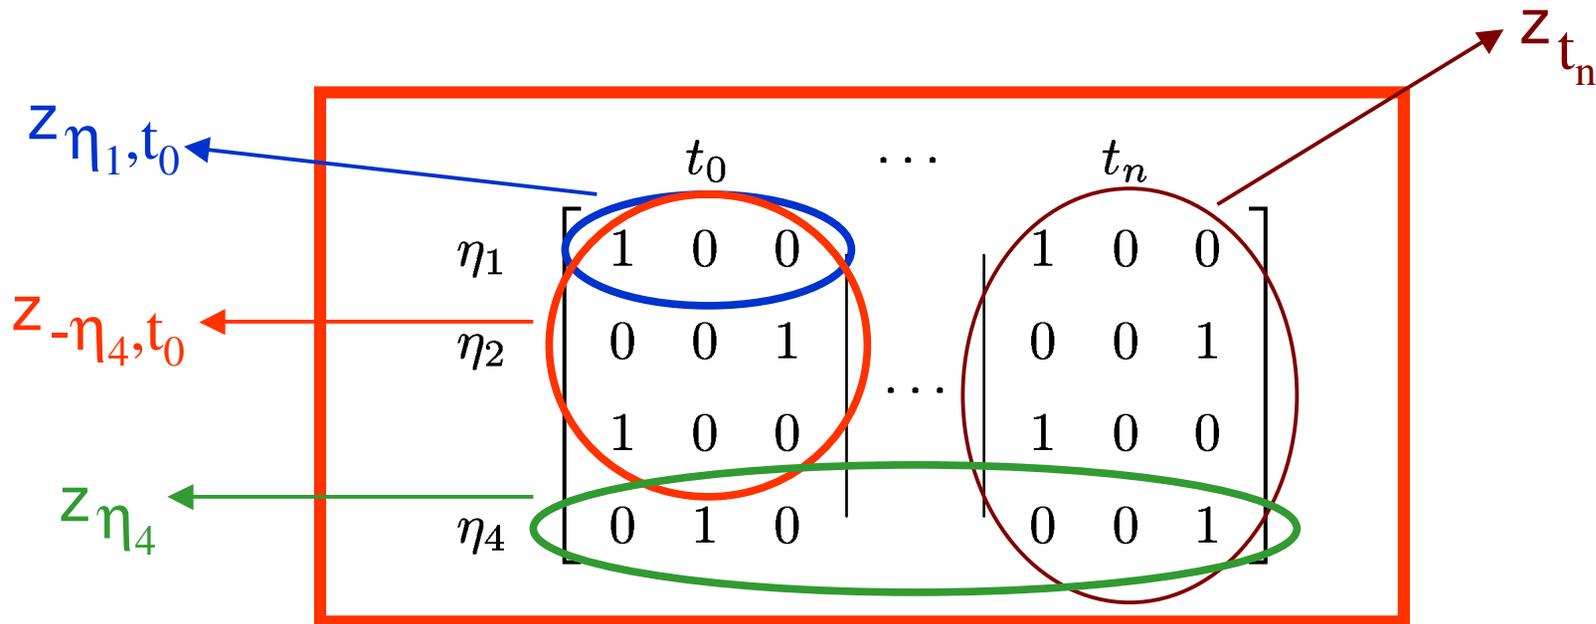
Clamping

Wonderful Life Utility



Nomenclature

- η : an agent
- z : state of all agents across all time
- $z_{\eta,t}$: state of agent η at time t
- $z_{-\eta,t}$: state of all agents other than η at time t





Aristocrat Utility

- One can solve for factored U with maximal learnability, i.e., a U with good term 2 and 3 in central equation:

$$\begin{aligned} AU_{\eta}(z) &\equiv G(z) - E[G(z) | z_{-\eta}] \\ &= G(z) - \sum_i p_i \cdot G(z_{-\eta}, CL_{\eta}^{s_i}) \end{aligned}$$

- Intuitively, AU reflects the difference between the actual G and the average G (averaged over all actions you could take).
- For simplicity, when evaluating AU here, we make the following approximation:

$$p_i(z_{\eta}) = \frac{1}{\text{Number of possible actions for } \eta}$$



Clamping

- **Clamping** parameter $CL_{\eta}^{\vec{v}}$: replace η 's state (taken to be unary vector) with constant vector \vec{v}
- Clamping creates a new “virtual” worldline
- In general \vec{v} need not be a “legal” state for η
- Example: four agents, three actions. Agent η_2 clamps to “average action” vector $\vec{a} = (.33 \ .33 \ .33)$:

$$\begin{array}{c} \eta_1 \\ \eta_2 \\ \eta_3 \\ \eta_4 \end{array} \begin{array}{c} z \\ \left[\begin{array}{ccc} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{array} \right] \end{array} \quad \begin{array}{c} \Rightarrow \\ \text{Clamp } \eta_2 \end{array} \quad \begin{array}{c} (z_{-\eta}, \vec{a}) \\ \left[\begin{array}{ccc} 1 & 0 & 0 \\ \oplus & \oplus & \oplus \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{array} \right] \end{array}$$



Wonderful Life Utility

- The Wonderful Life Utility (WLU) for η is given by:

$$WLU_{\eta}(z) \equiv G(z) - G(z_{-\eta}, CL_{\eta}^{\dot{v}})$$

- Clamping to “null” action ($\dot{v} = \vec{0}$) removes player from system (hence the name).
- Clamping to “average” action disturbs overall system minimally (can be viewed as approximation to AU).
- Theorem: WLU is factored regardless of \dot{v}
- Intuitively, WLU measures the impact of agent η on the world
 - Difference between world as it is, and world without η
 - Difference between world as it is, and world where η takes average action
- WLU is “virtual” operation. System is **not** re-evolved.

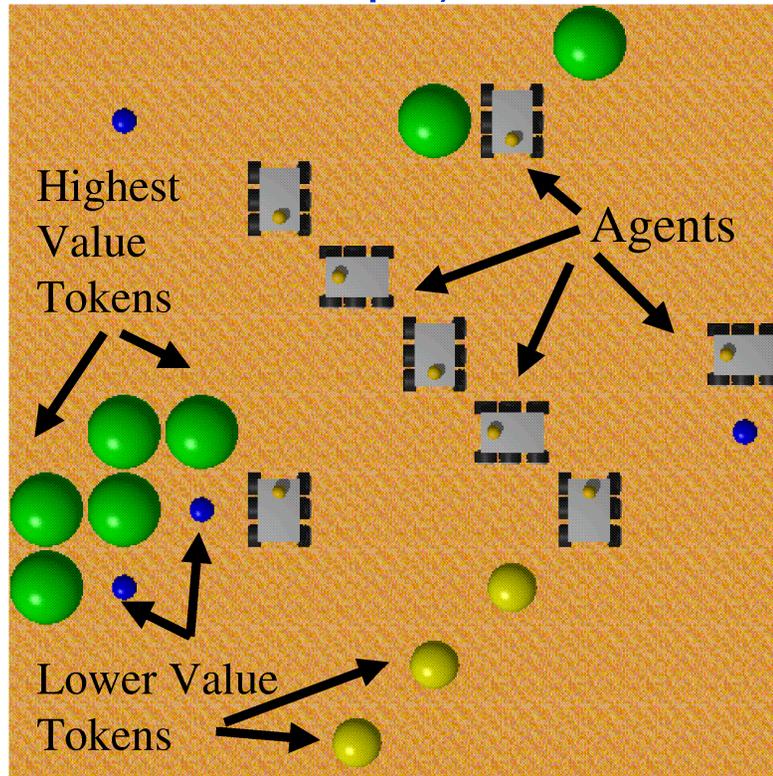
DETAILS FOR THE ROVER PROBLEM:

Formulation
World Utility
Payoff Utilities
Results



Collectives of Rovers

- Design a collective of autonomous agents to gather scientific information (e.g., rovers on Mars, submersibles under Europa)



- Some areas have more valuable information than others
- World Utility: Total importance weighted information collected
- Both the individual rovers and the collective need to be flexible so they can adapt to new circumstances
- Collective-based payoff utilities result in better performance than more “natural” approaches



Rover Problem: World Utility

- Token value function:

$$V(L, \Theta) = \sum_{x,y} \Theta_{x,y} \min(1, L_{x,y})$$

- L : Location Matrix for all agents
- L_{η} : Location Matrix agent η
- $L_{\eta,t}^a$: Location Matrix of agent η at time t , had it taken action a at $t-1$
- Θ : Initial token configuration

- World Utility :

$$G(z) = V(L, \Theta)$$

- Note: Agents' payoff utilities reduce to figuring out what “ L ” to use.





Rover Problem: Payoff Utilities

- Selfish Utility :

$$SU_{\eta}(z) = V(L_{\eta}, \Theta)$$

- Team Game Utility :

$$TG_{\eta}(z) = V(L, \Theta)$$

- Collectives-Based Utility (theoretical):

$$AU_{\eta}(z) = G(z) - \sum_{a \in A_{\eta}} p_a^r V(L_{\wedge \eta} + L_{\eta}^a, \Theta)$$

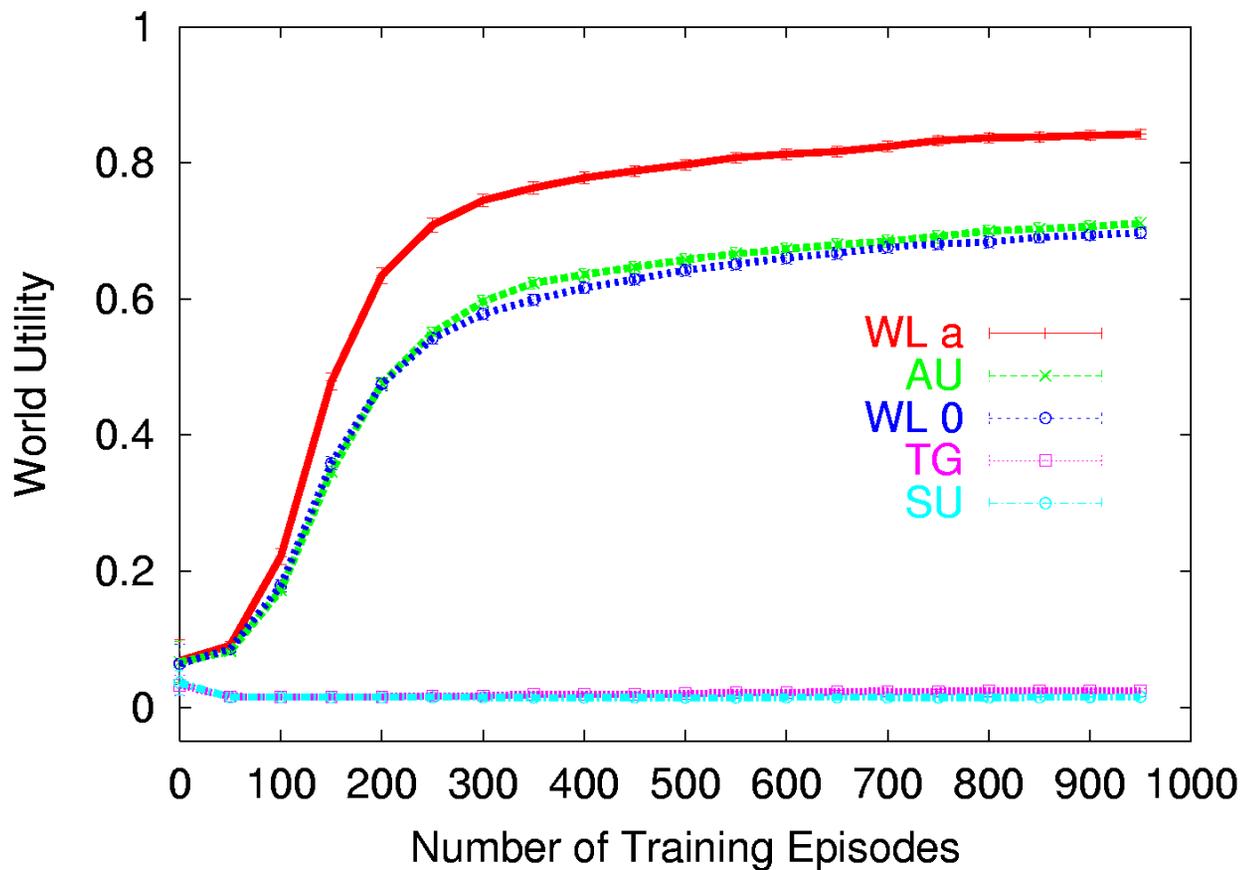
- Collectives-Based Utility

$$WLU_{\eta}^a(z) = G(z) - V(L_{\wedge \eta} + \sum_{a \in A_{\eta}} p_a^r L_{\eta}^a, \Theta)$$

$$WLU_{\eta}^{\emptyset}(z) = G(z) - V(L_{\wedge \eta}, \Theta)$$



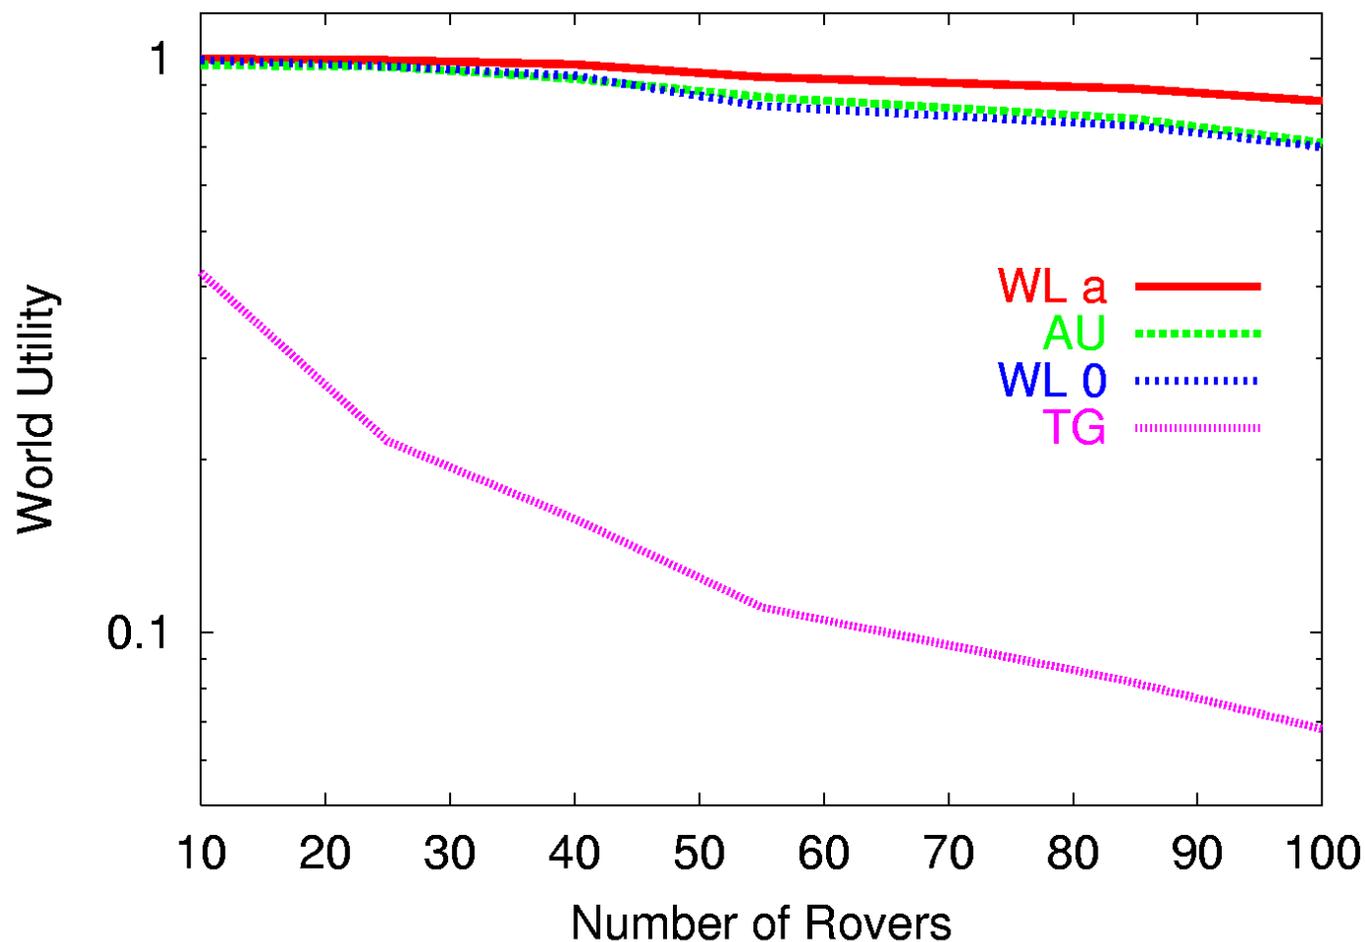
Rover Problem: Utility Comparison



100 rovers on a 32x32 grid



Rover Problem: Scaling Properties



DETAILS FOR THE AUTONOMOUS DEFECTS PROBLEM

Formulation

World Utility

Results



Autonomous Defects Problem

- Given a collection of faulty devices, how to choose the subset of those devices that, when combined with each other, gives optimal performance (Johnson & Challet).

$$G(z) = \frac{\sum_{j=1}^N n_j a_j}{\sum_{k=1}^N n_k}$$

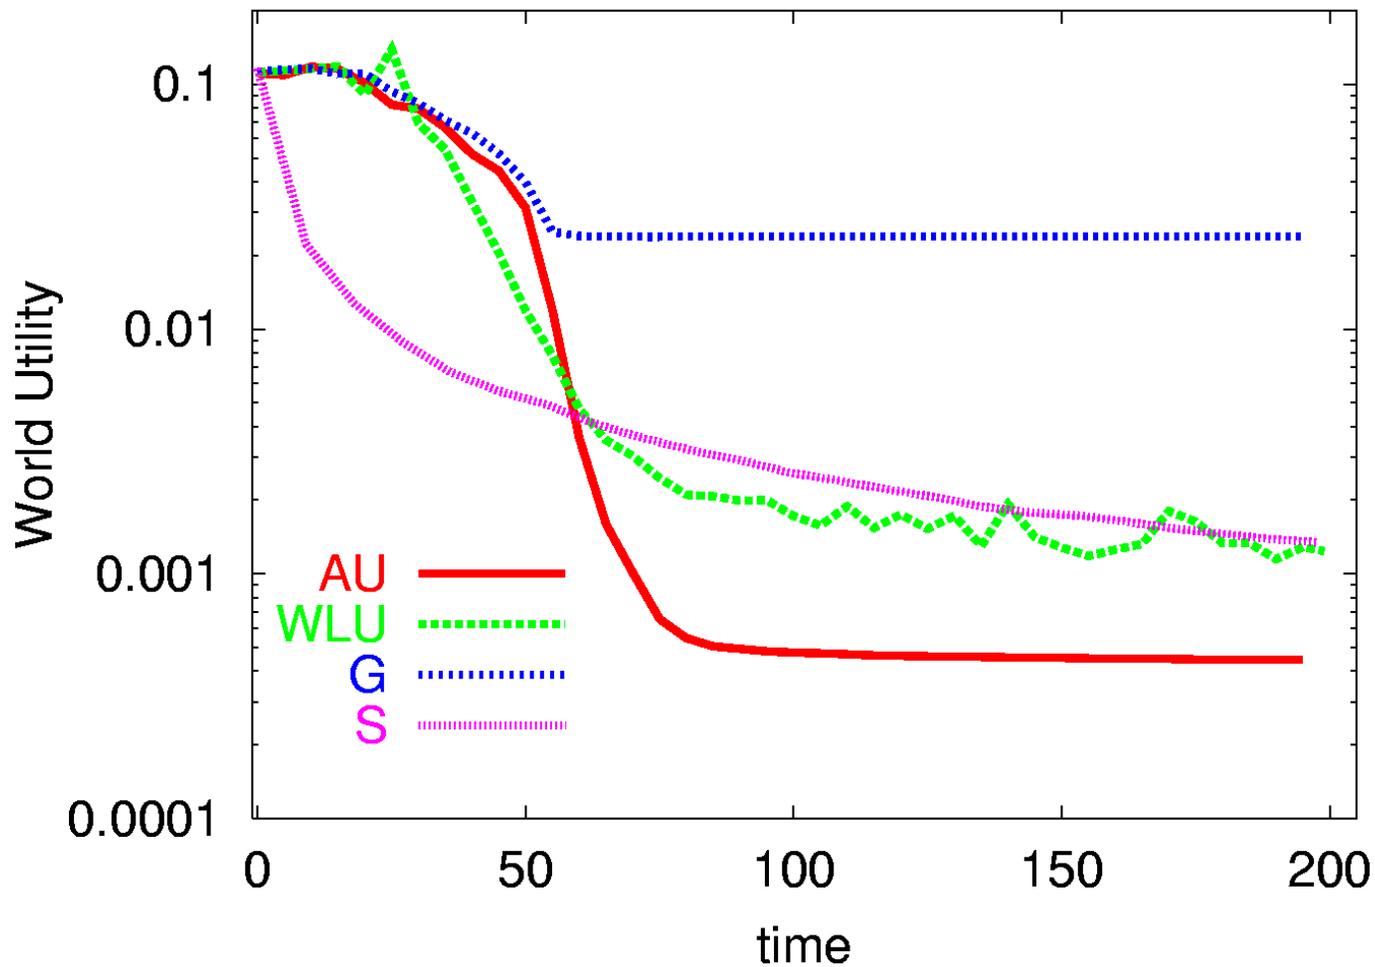
a_j : distortion of component j

n_k : action of agent k ($n_k = 0 ; 1$)

- Collective approach: Identify each agent with a component.
- Question: what utility should each agent try to maximize?

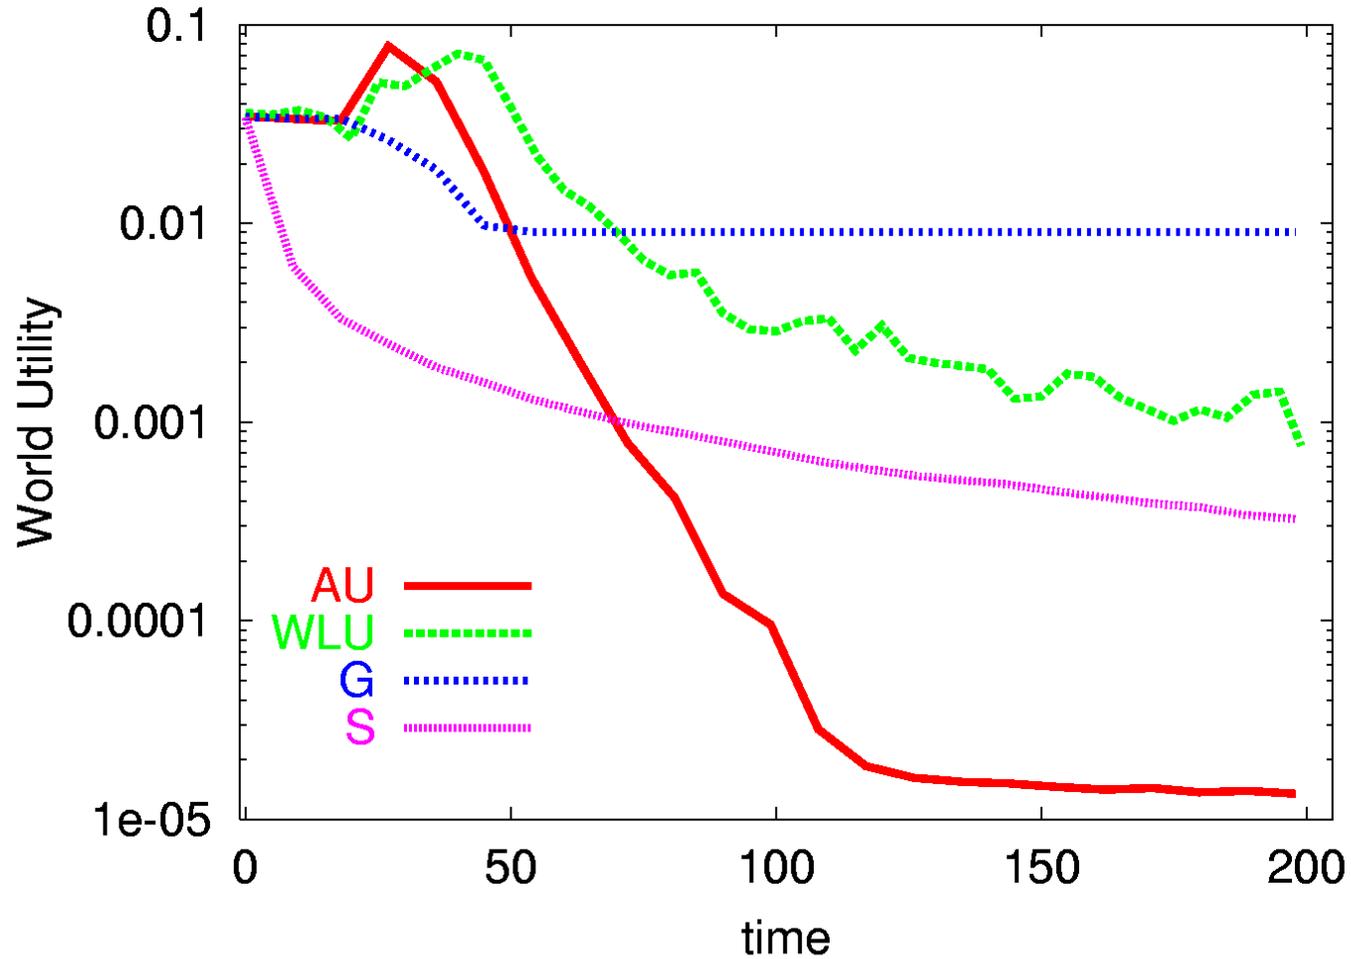


Autonomous Defects Problem (N=100)



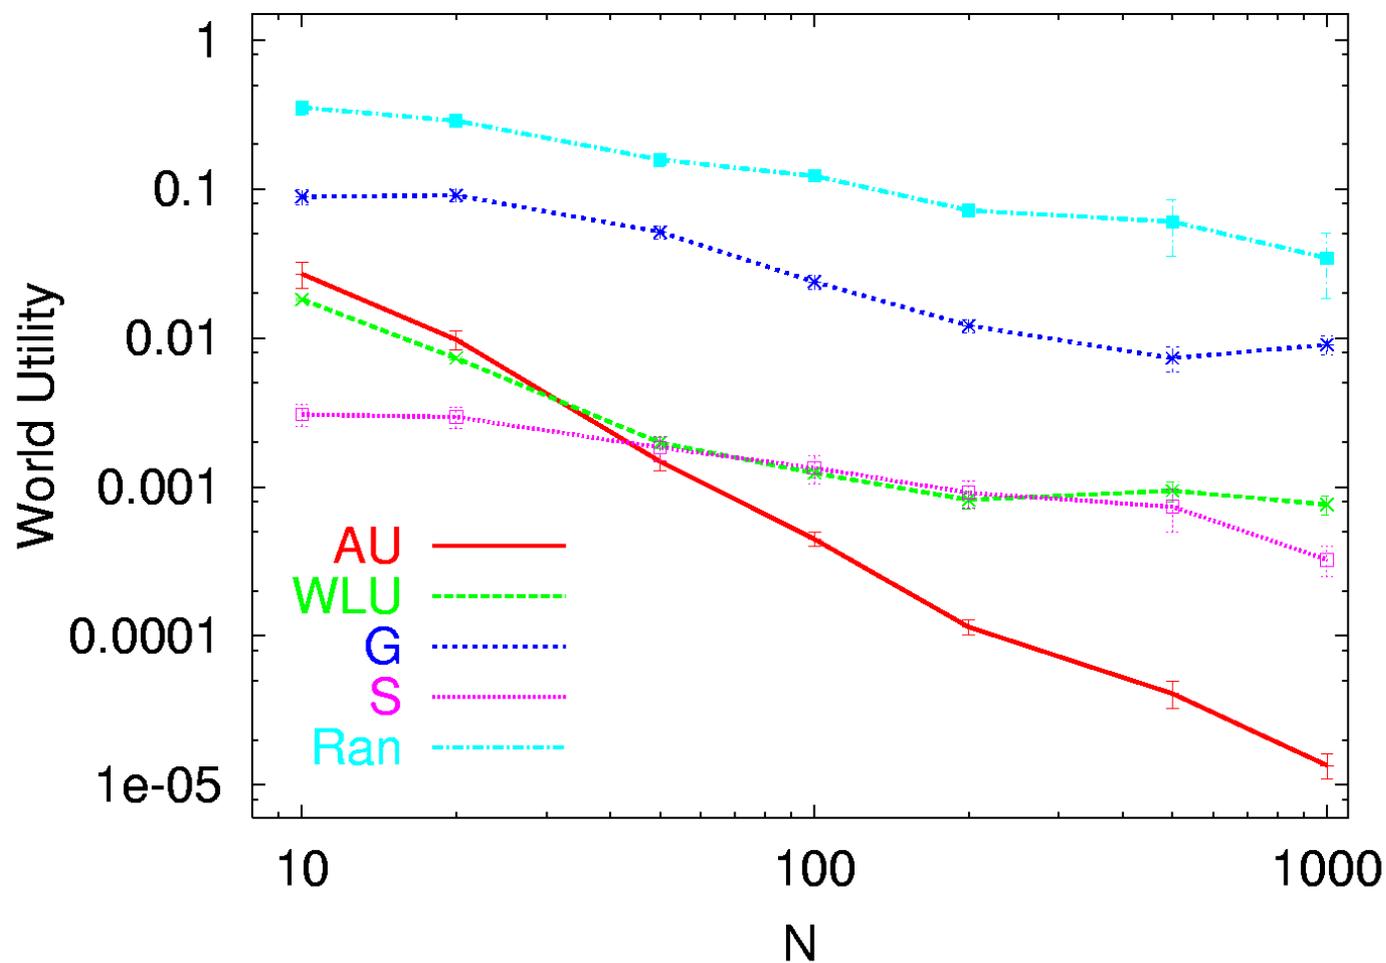


Autonomous Defects Problem (N=1000)





Autonomous Defects Problem: Scaling





Constellation of Satellites

- Problem:
 - A set of satellites receives data faster than they can download (eg., in orbit around Earth, or for that matter Mars)
 - Cannot be centrally controlled (size, and communication delays)
- Approach:
 - Adaptively route data to minimize importance weighted data loss
 - Investigated “fooling” a baseline algorithm by introducing “ghost” traffic
 - Agents set ghost traffic using theory of collectives