Taxonomy

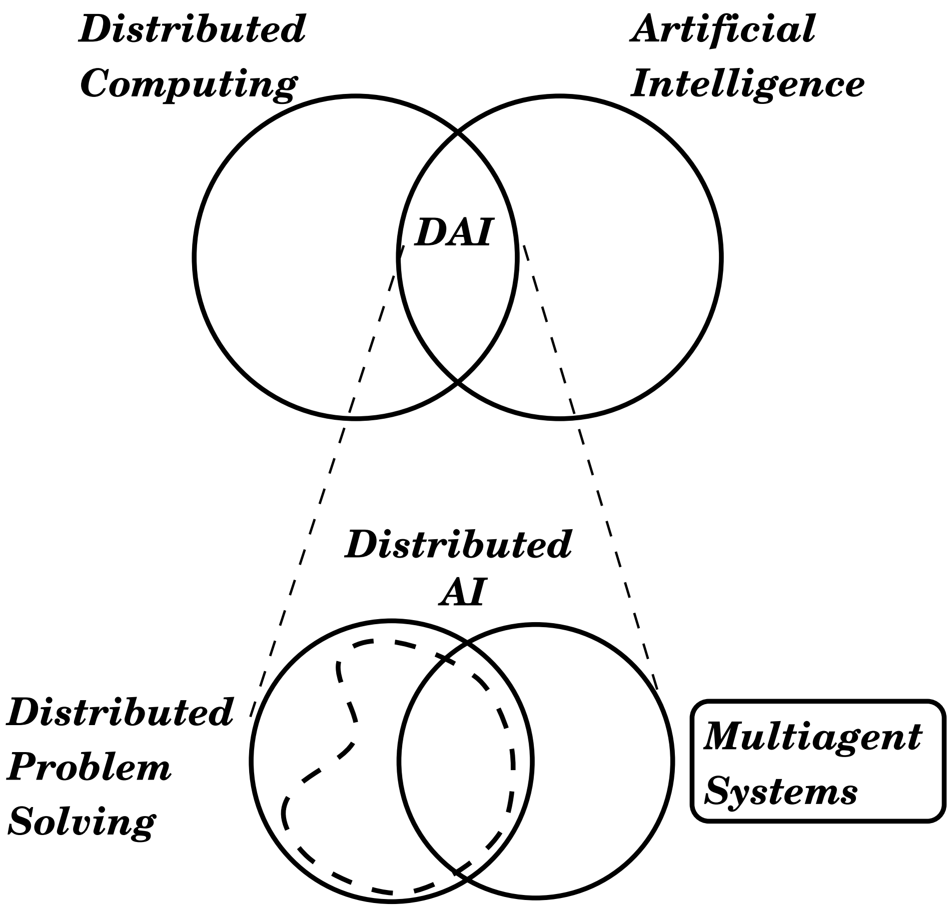

Distributed Artificial Intelligence (DAI) is concerned with the study and design of systems consisting of several interacting entities which are logically and often spatially distributed and in some sense can be called autonomous and intelligent. *Gerhard Weiss*

Definitions

Distributed Computing : Several processors share data, but not control. Focus on low-level parallelization or synchronization issues.

Distributed AI : Control as well as data is distributed. Focus on problem solving, communication, and coordination.

Distributed Problem Solving : Task decomposition (task sharing) and/or solution synthesis (result sharing): information management.

Multiagent Systems : Behavior coordination or behavior management.

- No necessary guarantees about other agents.
- Individual behaviors typically simple relative to interaction issues.

Why Multiagent Systems?

- Some domains require it. (Hospital scheduling)
- Parallelism.
- Robustness.
- Scalability
- Simpler programming.
- "Intelligence is deeply and inevitably coupled with interaction." – *Gerhard Weiss*

The Predator-Prey ("Pursuit") Domain

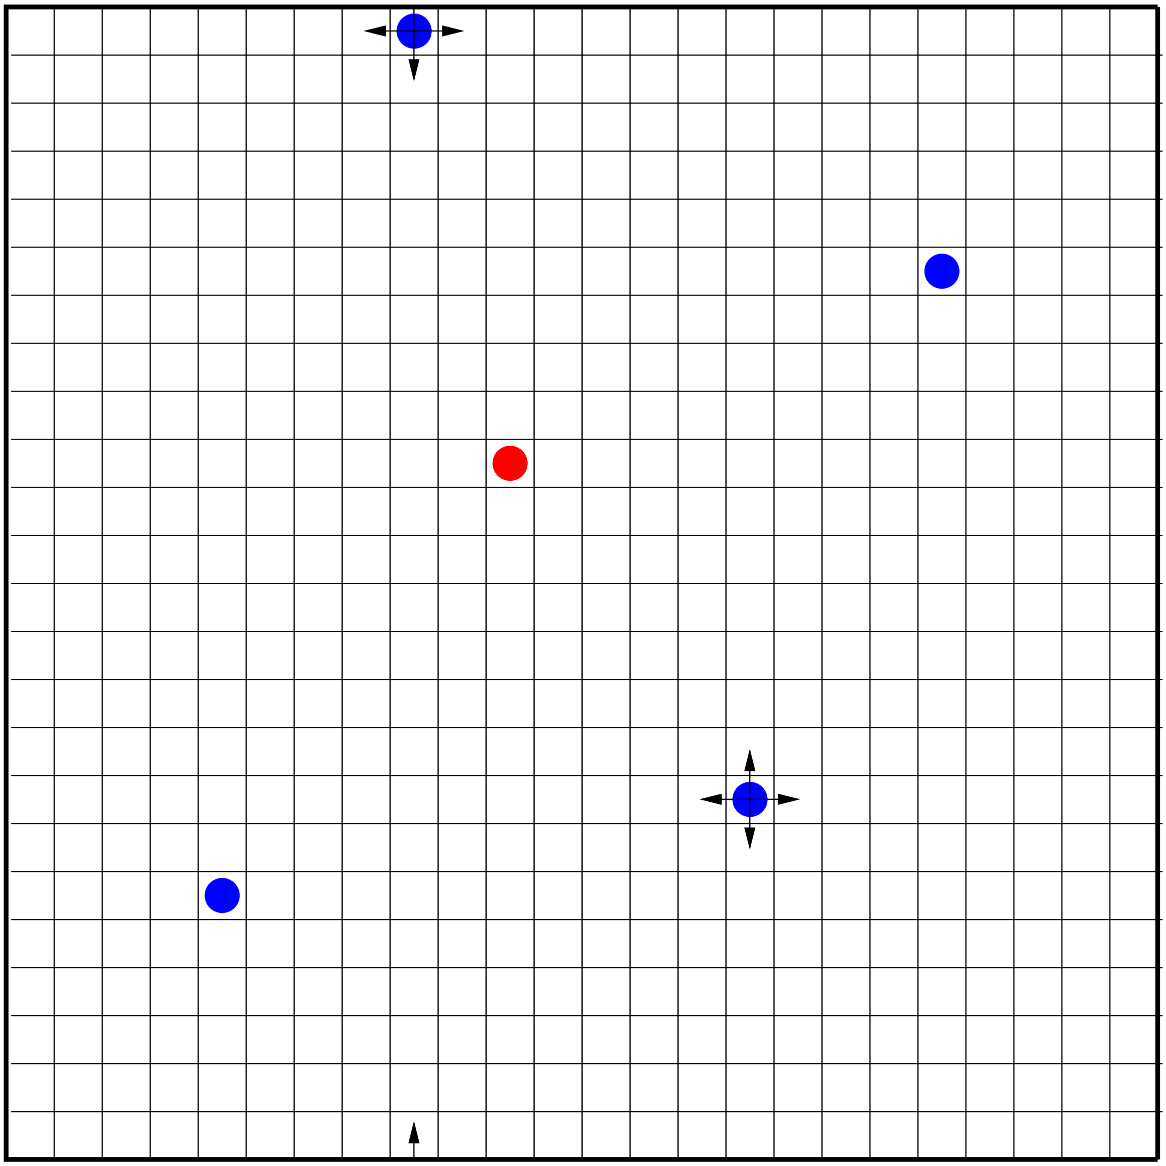

Orthogonal Game in a Toroidal World

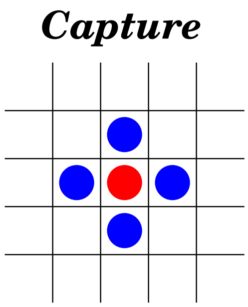

- *Predators see each other*
- *Predators can communicate*
- *Prey moves randomly*
- *Prey stays put 10% of time*
- *Simultaneous movements*

Single Agent Pursuit

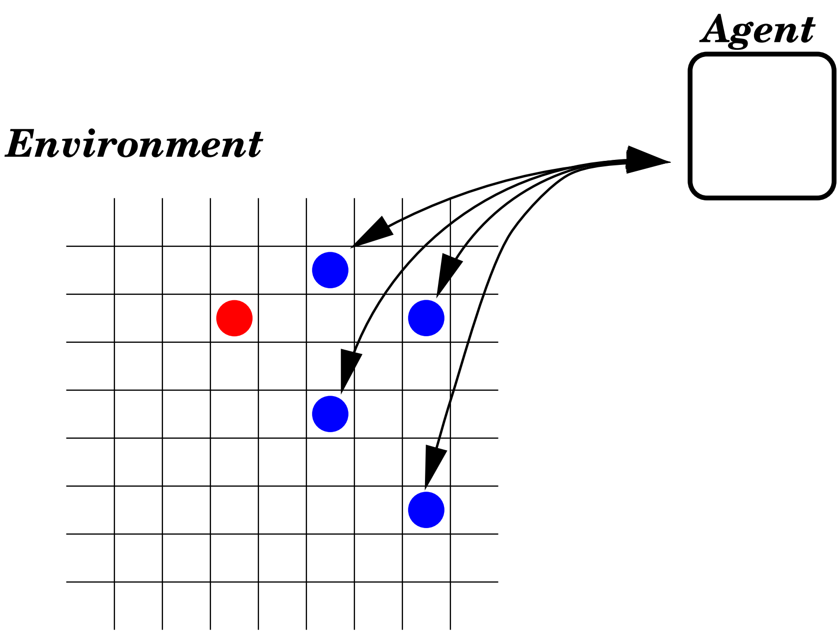

- One agent controls all predators.
- Prey is considered part of the environment.

Single Agent Systems

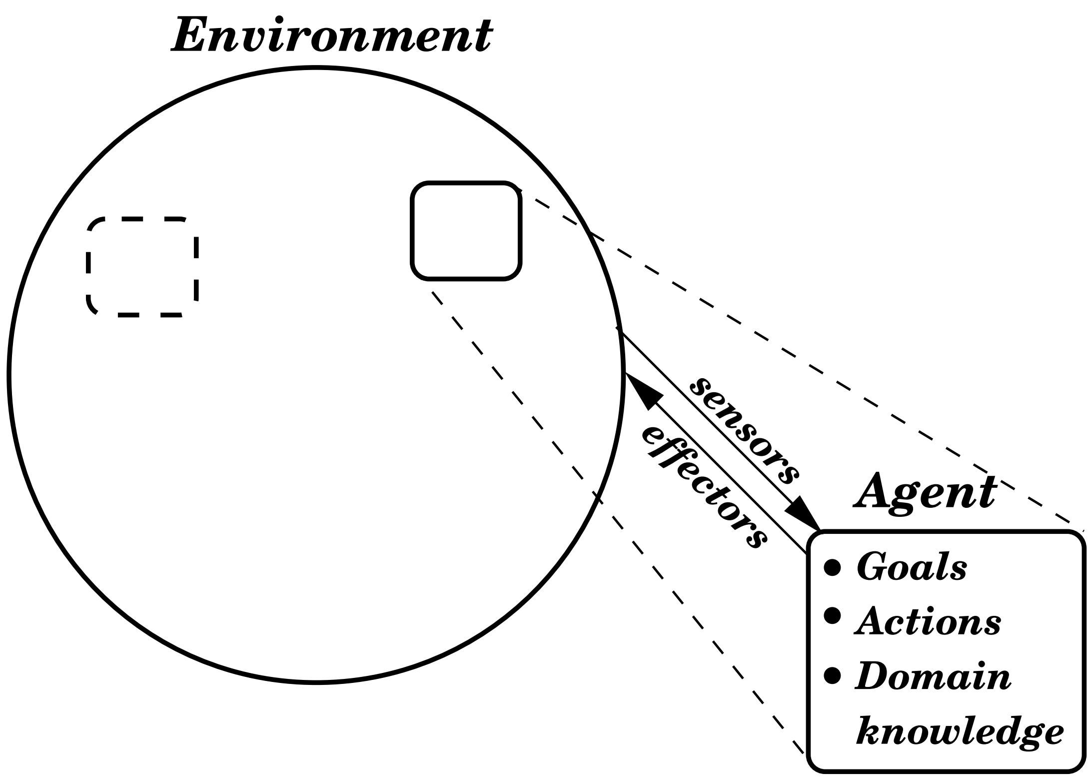

- Agent models itself, the environment, and their interactions.
- If other agents exist, they're considered environment.

Multiagent Systems

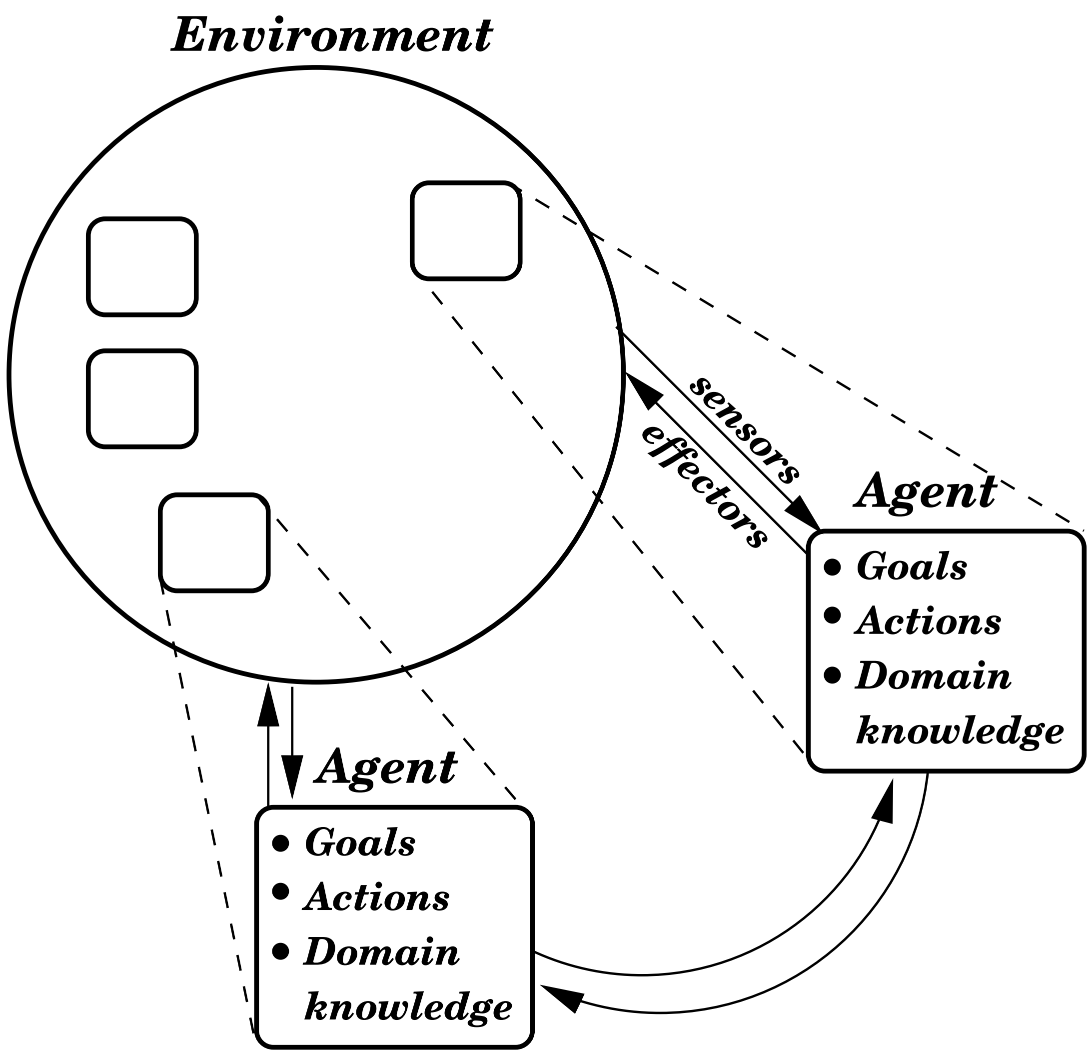

- Agents model each others' goals and actions.
- Possible direct interaction.
- Dimensions: degree of agent **heterogeneity**; degree of **collaboration**; **communication** capabilities.

Homogeneous Multiagent Pursuit

- One **identical** agent per predator.
- Possibly limited information about other agents' internal states.

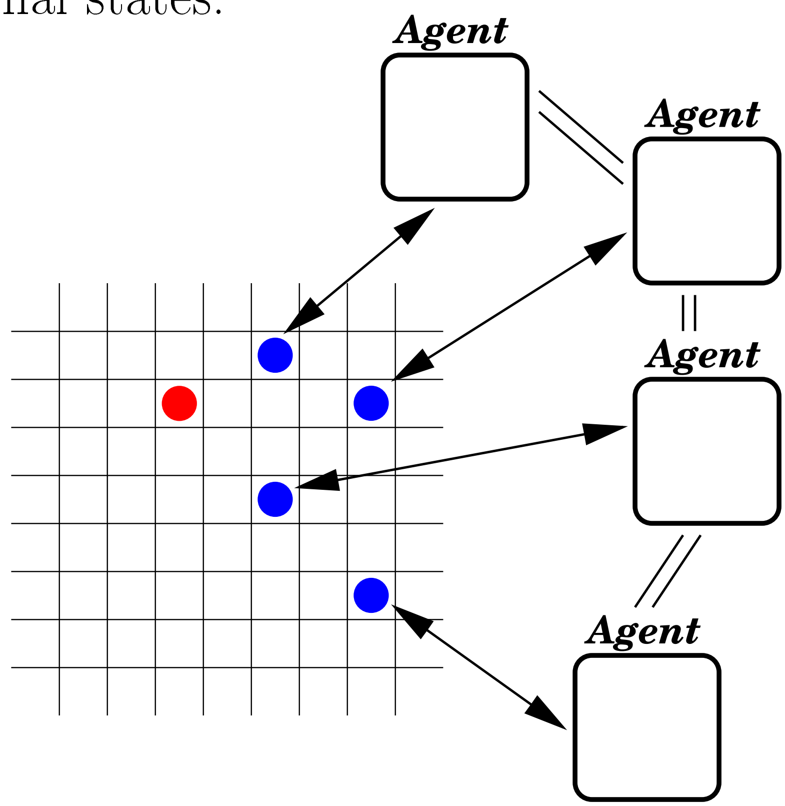

- Each agent to closest capture position. *Stephens and Merx*
- Predict other agents' moves using **Recursive Modeling**. *Vidal and Durfee*
- **Game theory**: each agent maximizing own utility leads to coordination. *Levy*
- Claim that game is easily solved with greedy decisions (diagonal, hexagonal games): attract to prey, repel other predators. *Korf*

Homogeneous Agents

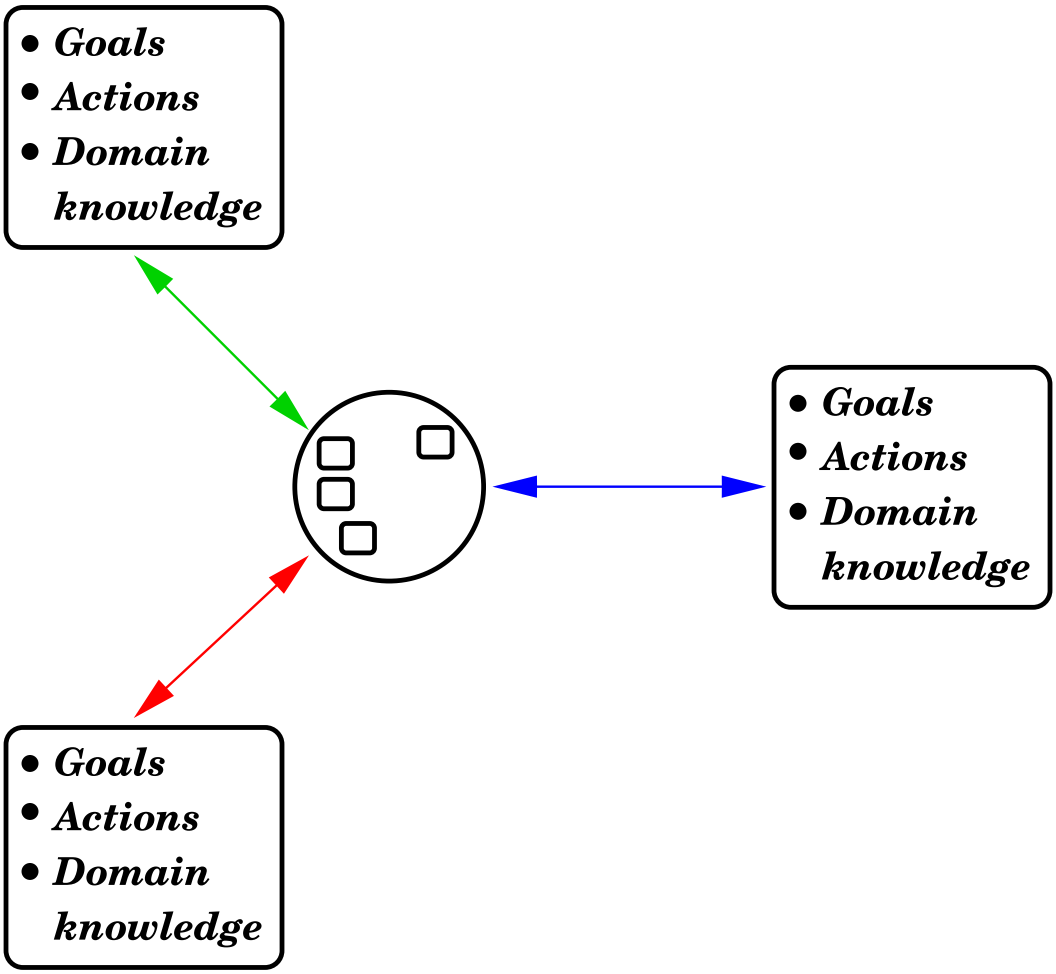

- Only sensors and effectors differ.
- Necessary condition.
- For now assume no direct interaction.

Heterogeneous Multiagent Pursuit

- Prey also modeled as an agent.
- Goals and actions differ among agents.

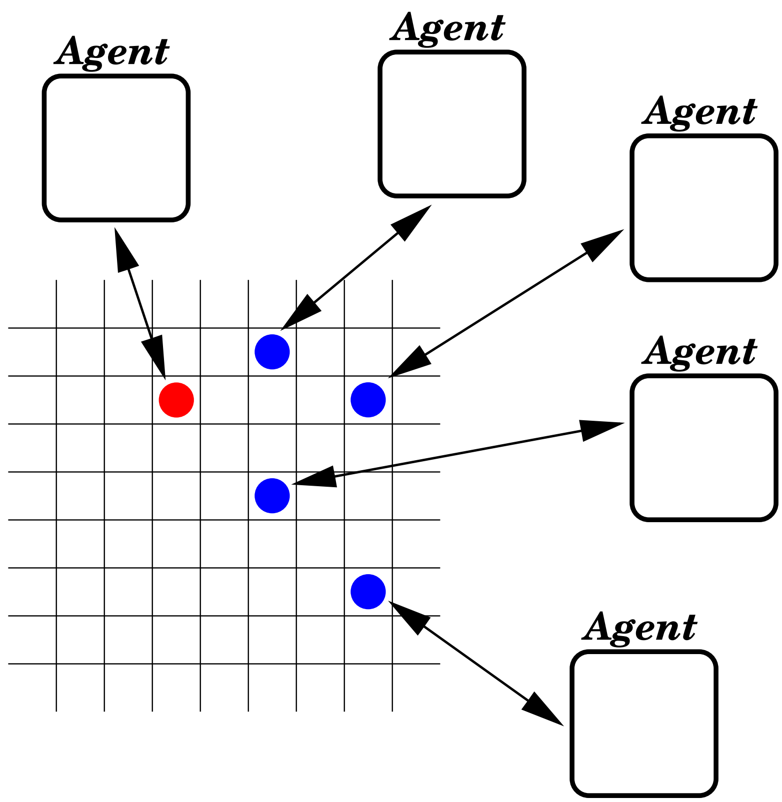

- Evolve teams of predators: implicit communication. *Haynes and Sen*
- Model others by set of default rules, then store negative cases and update behaviors (no co-occupation). *Haynes and Sen*
- **Evolve prey behaviors** (linear always succeeds?). *Haynes and Sen*

Heterogeneous Agents

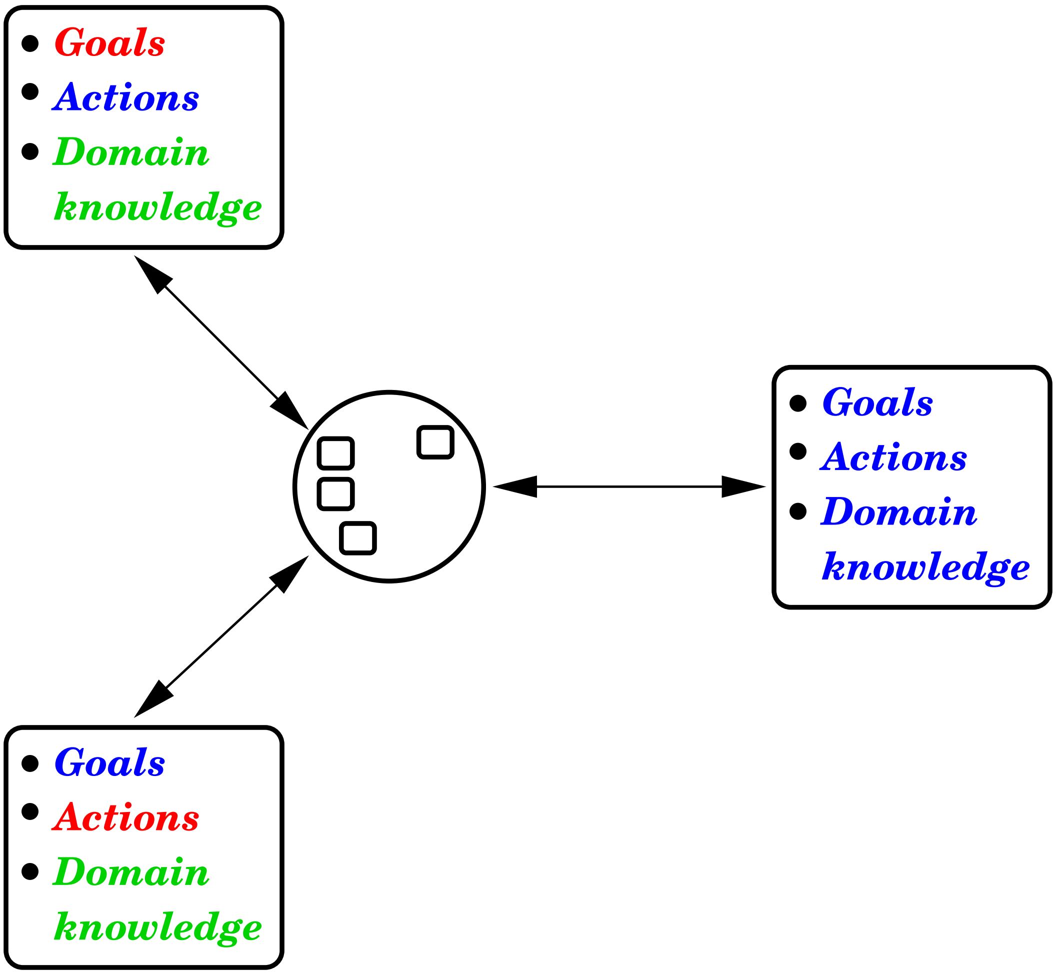

- Goals, actions, and/or domain knowledge may differ.
- Still assume no direct interaction.

Communicating Multiagent Pursuit

- Predators can communicate.

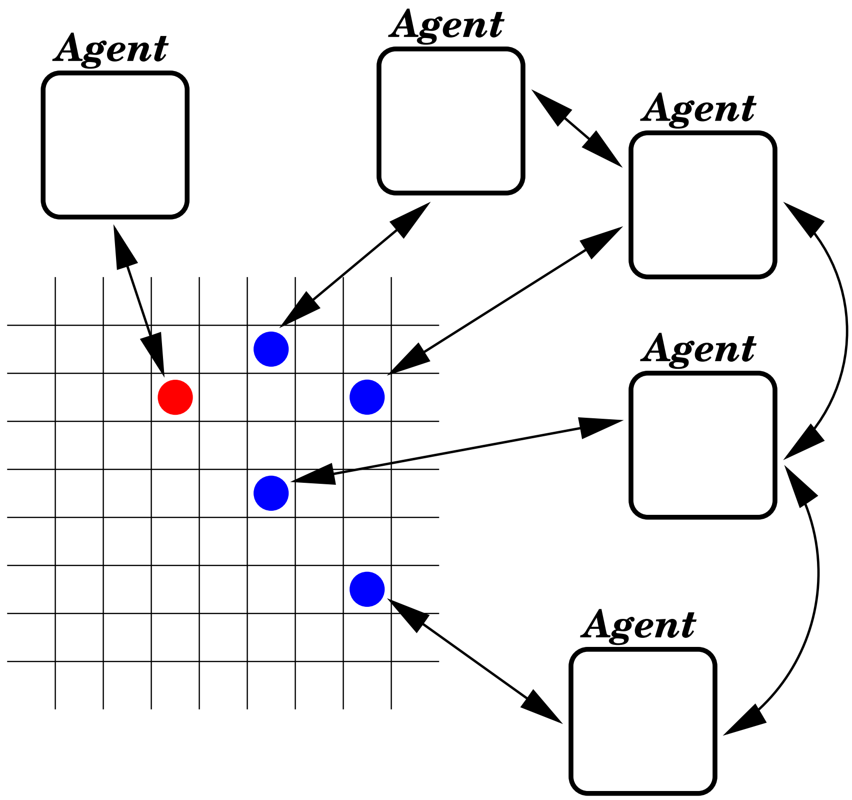

- Predators use Q-learning and share either sensation, episodes, or policies (predators, scouts, limited view). Share sensation for joint task. *Tan*
- Choose capture position to approach one at a time: farthest first. *Stephens and Merx*
- One agent controls all. *Stephens and Merx*

- Exchange data; exchange data and goals; or give up control: tradeoff between communication costs and better decisions. *Benda et al.*
- Start autonomously, then communicate, then negotiate, then control. Cost increases as freedom decreases. *Osawa*

In the extreme, we get back to the single agent scenario

Communicating Agents

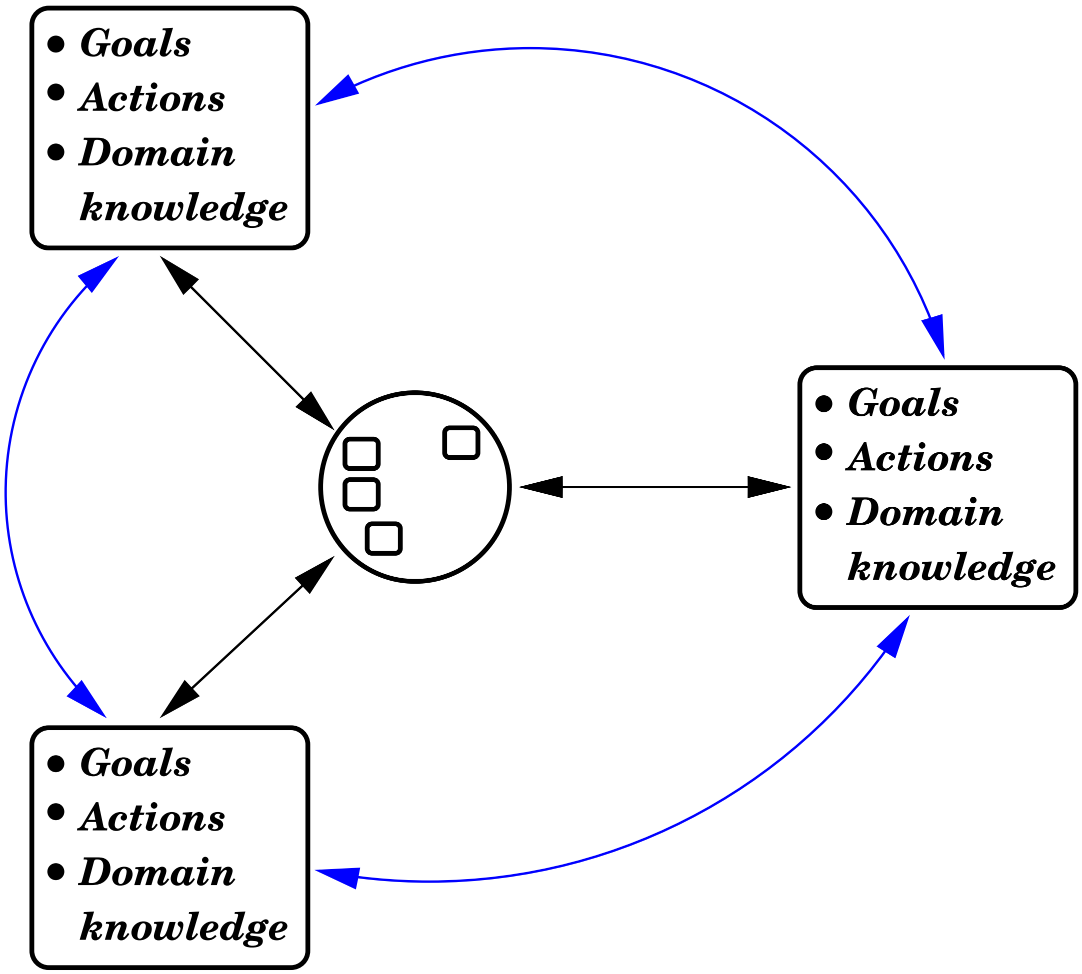

- Information can be transmitted directly among agents.
- Blackboard or point-to-point communication.