CMRoboBits: Multi-Robot Planning, Coordination

Manuela Veloso
Scott Lenser
Thanks to Douglas Vail

15-491
December 1st, 2003

Dynamic Multi-Robot Environments

Robots need:

- To perceive and build a model of the environment.
- To achieve goals and respond in real time.
- To integrate deliberative planning and reactive execution.

Robot soccer: A concrete multirobot research challenge.

Sony AIBO Robots

Fully autonomous: vision, computation, motion

2002: wireless communication

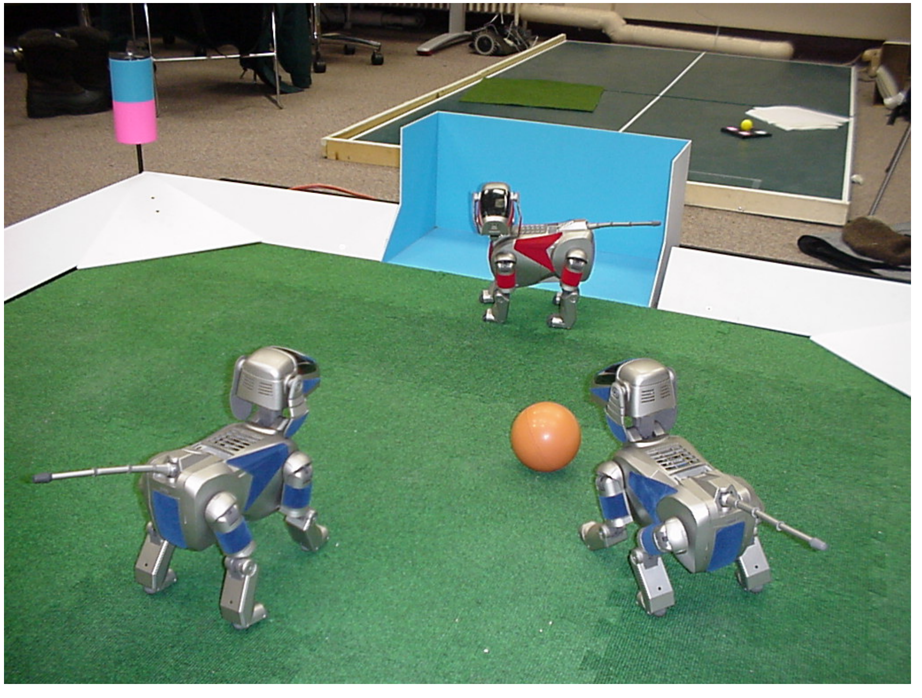

The Physical Setup for the Robot

Single Robot Performance

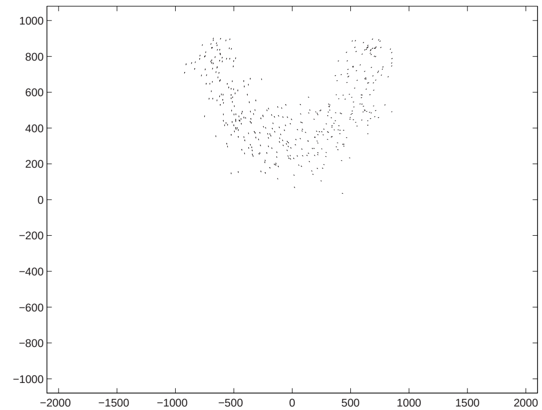

Color thresholding, object recognition

CMVision
Bruce, Lenser

Sensor-resetting localization

SRL
Lenser

Walking, kicking

4-Legged Motion
Chernova, Bruce

Robustness in multi-robot environments.

Multiple Robots

- How are the robots a **team** of robots?
- 1998-2001: No communication among robots
- Apriori fixed roles - goalie, attackers and multi-fidelity behaviors, i.e., multiple behaviors as a function of confidence in world model
- Still interesting *emerging* teamwork
- 2002: Wireless communication

Team Experience in Other Platforms

- Simulation - fully distributed, communication
- Small-size - global view of the world

Communication and World Model

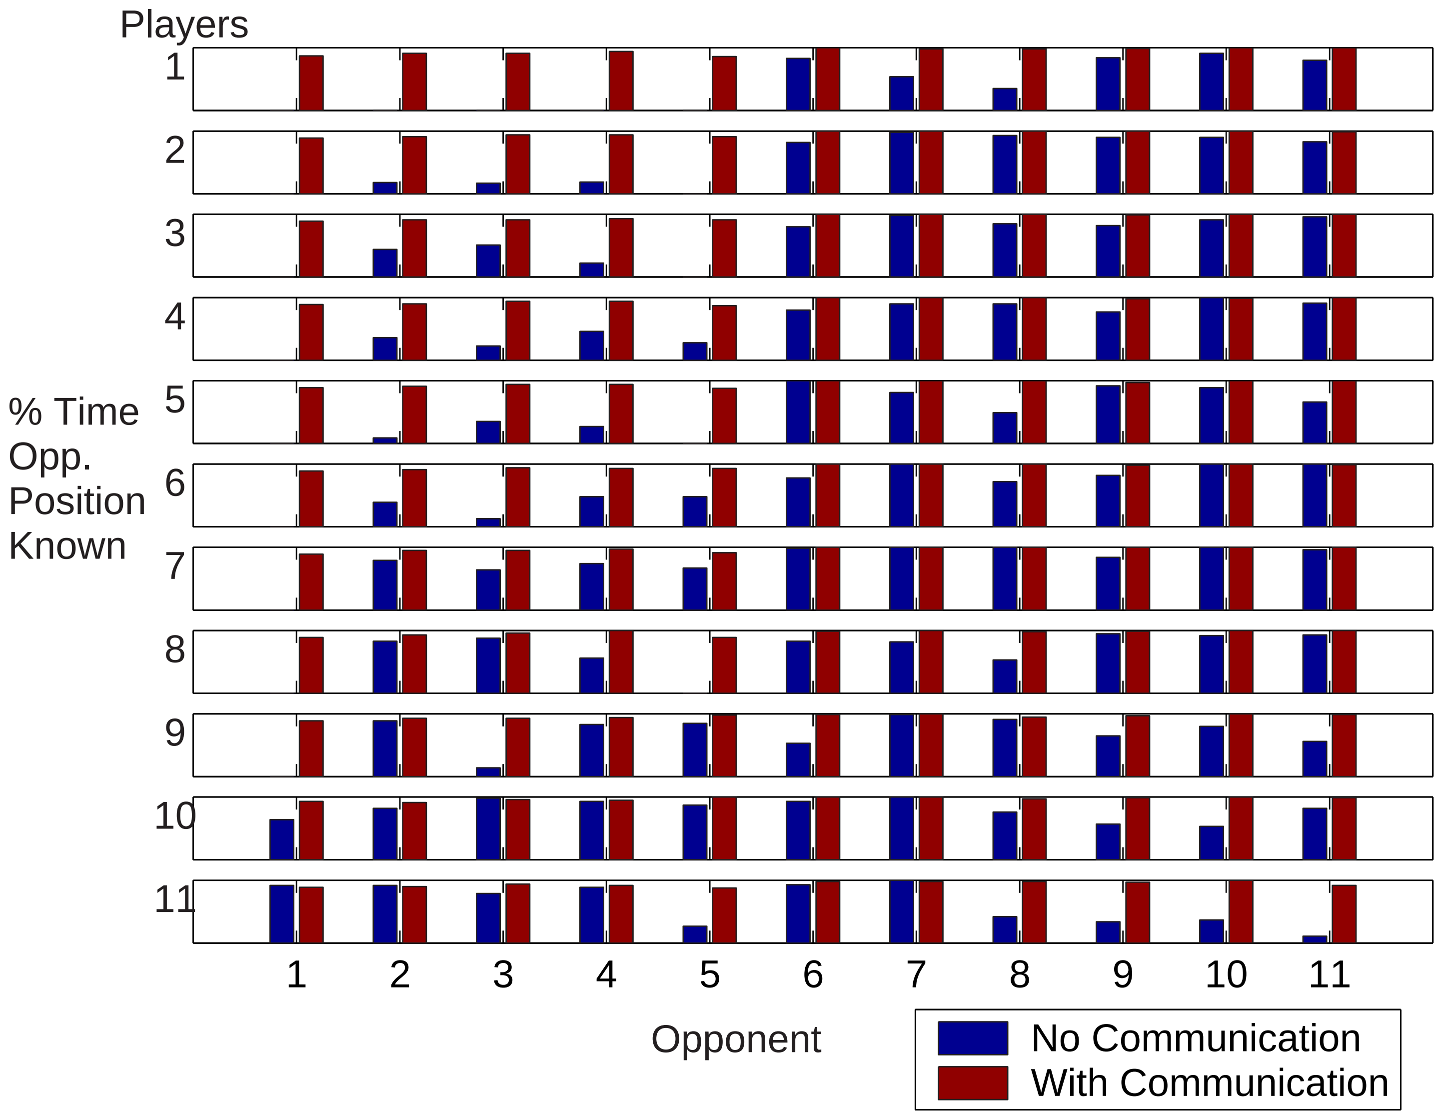

Communication to Share Information

- Broadcast positions and confidences:
  - Robot's: x_r, y_r, σ_r and Ball's: x_b, y_b, σ_b
- Build *individual* and *shared* world models
  - Individual - from own perception, memory, and prediction
  - Shared - from teammate robots
- Shared world model is accessed only on demand
  - Rely on individual model up to some confidence
  - Use shared if low confidence in individual

Roth, Vail, & Veloso, 2003, under submission

Multi-Robot Coordination

- Given a world model, coordination involves two phases:
  - Role assignment
  - Positioning

Ordered Role Assignment

Given:

- An ordered list of n roles: $r_1, \ldots, r_n$
  - Primary attacker, defensive supporter, offensive supporter
- A list of n robots: $R_1, \ldots, R_n$
- A value for a role, function of a robot: $\text{Bid}_{r_i}(R_j)$
- An initially empty assignment: $\text{assigned}(R_i) = 0, \forall i.$

Compute a role assignment:

1. For $i = 1, \ldots, n$
2. $\text{assigned}(R_j) = r_i$, s.t.

   $\text{Bid}_{r_i}(R_j) > \text{Bid}_{r_i}(R_k), \forall k \neq j \wedge \text{assigned}(R_k) = 0$

Role Assignment Bidding

- Primary attacker:

$$\text{Bid}_{\text{attacker}}(R) = \underbrace{\frac{\theta_{goal}(R)}{\pi}}_{angular\ component} + \underbrace{(1 - min(1, d_{ball}(R)))}_{distance\ component} \quad (1)$$

- Offensive supporter

- Defensive supporter

Role Assignment Issues

- Local-robot evaluation of multi-robot bids
  - Information is communicated for multiple purposes
- Latency of communication
  - Hysteresis - robots hold roles
- Cost and frequency of communication
  - Ongoing open questions, on demand, differential

Roles for Multi-Robot Coordination


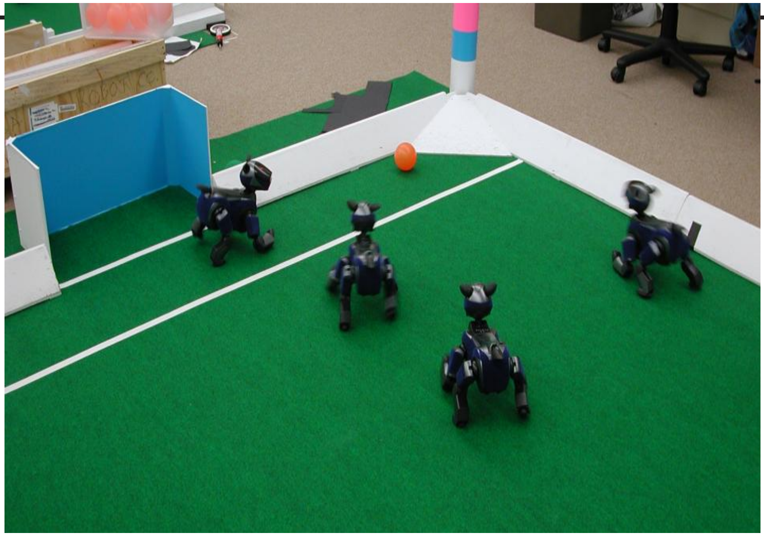

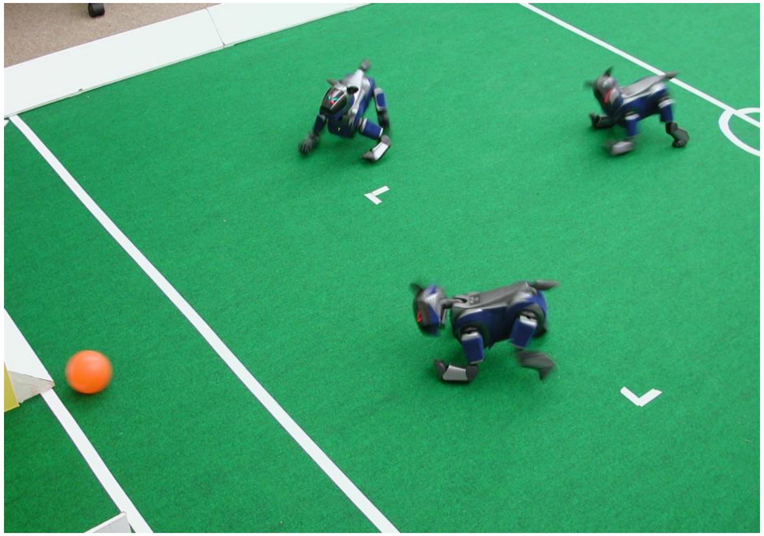

Potential Functions for Positioning

- Potential functions as a sum of several linear *components* of two different types:
  - Constraints about the task, e.g., do not block shots
  - Dynamic obstacle information, e.g., repulsion from walls and other robots

- Different components are used in different roles.

Dynamic Positioning

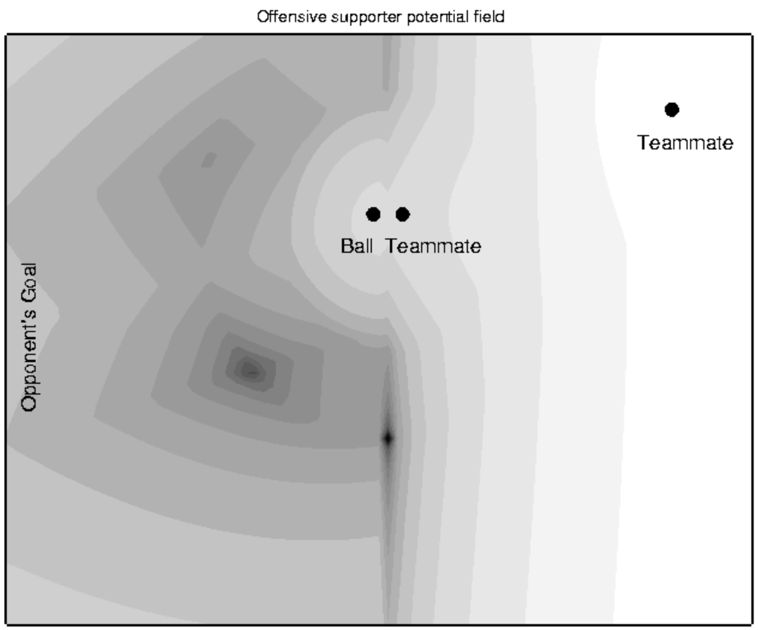

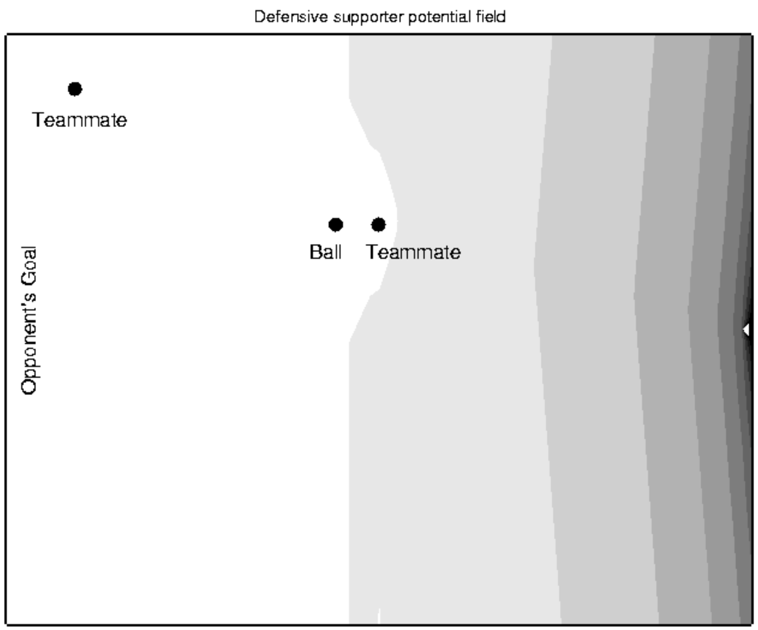

Offensive Supporter

Offensive Defender

Examples of Potential Components

- $P_{wall} = max(0, c_1 - k_1 \cdot d_{wall})$
- $P_{ball} = \|c_2 - d_{ball}\| \cdot k_2$
- $P_{teammate} = max(0, c_3 - k_3 \cdot d_{teammate})$
- $P_{forward\, bias} = max(0, k_4 \cdot d_{behind\, ball})$
- $P_{defensive\, bias} = k_5 \cdot d_{from\, goalline}$
- $P_{ball\, corridor} = \|c_6 - d_{shot\, path}\| \cdot k_6$
- $P_{side\, bias} = max(0, k_8 \cdot offset_{robot} \cdot \frac{offset_{ball}}{\frac{1}{2} \cdot width_{field}})$

Ball Potential Component

$$P_{ball} = \|c_2 - d_{ball}\| \cdot k_2 \quad (2)$$

- It guides the offensive supporter to a position that is an equilibrium distance, c_2, away from the ball. The potential increases linearly with a slope of k_2 as the robot moves away from the equilibrium distance.

Games Scores in Robocup-2002

Opponent	Final score (CMPack'02:Opponent)
ARAIBO, Japan	5:1
German Team, Germany	3:1
SPQR, Italy	7:0
Team Sweden, Sweden	9:0
Robomutts++, Australia	4:0
rUNSWift, Australia	3:3, penalty kicks 2:1

Carnegie Mellon's CMPack'02 was World Champion.

Controlled Experiments

- Time to score: Three uncoordinated robots; three coordinated robots; one single robot

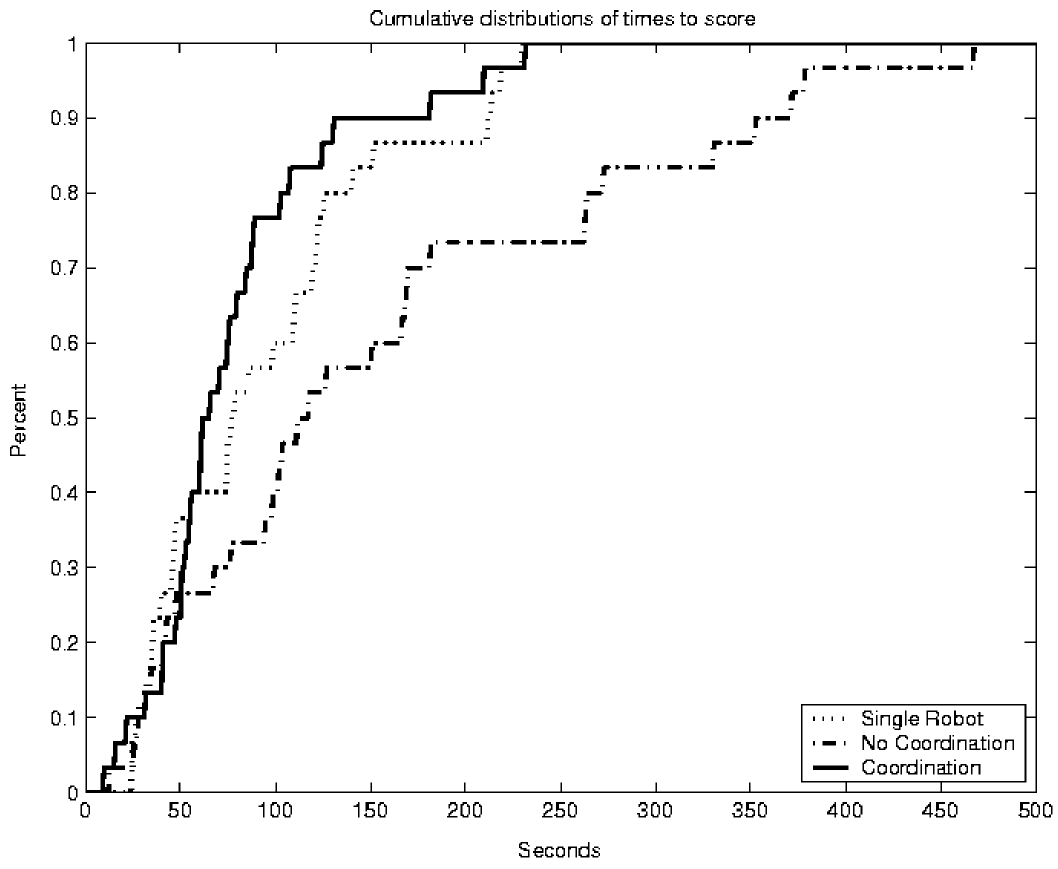

Ongoing Research

- Multi-robot behaviors as playbook
- Multi-robot learning
- Coaching
- Robust role assignment
- Communication control

Playbook for Team Plans

Team Playbook - **Locker-Room Agreements**

- Pre-defined protocols, actions, plans

- Triggered by world features visible to all

- CMDragons'02 - adaptation to the opponent:
  - definition of plays as *team plans*
  - multiple alternative plays in playbook
  - statistical gathering of success/failure/reward
  - weighted play selection

- Adaptation without explicit opponent modeling!

Multiagent Learning

The goal: A rational and convergent learning algorithm.

WoLF solution: A variable learning rate.

WoLF: Win or Learn Fast!

Michael Bowling PhD thesis, May 2003

A variable learning rate can achieve convergence without sacrificing rationality.

Results – Rock-Paper-Scissors

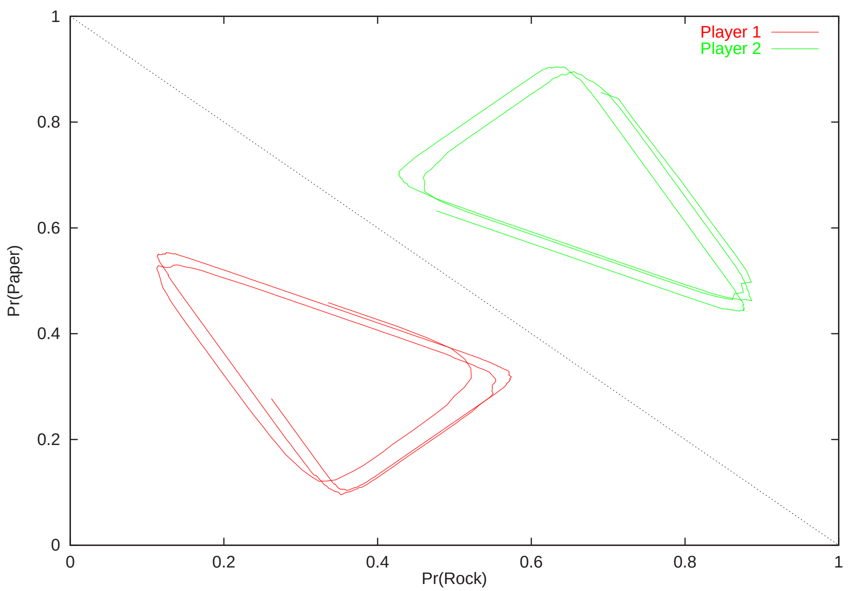

PHC

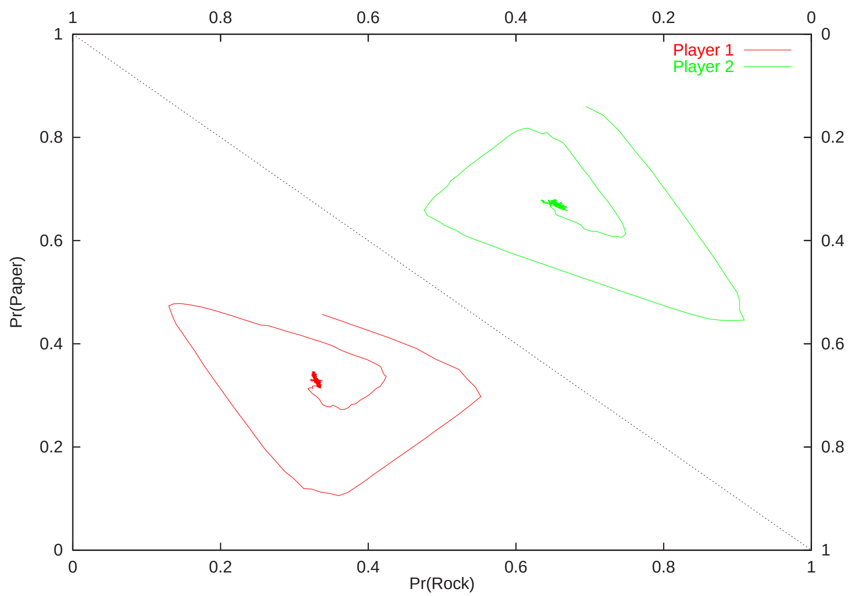

WoLF PHC

Learning in the Presence of Limitations

- Robots have limitations.
- WoLF converges.

Video

Coaching - Generation of Advice

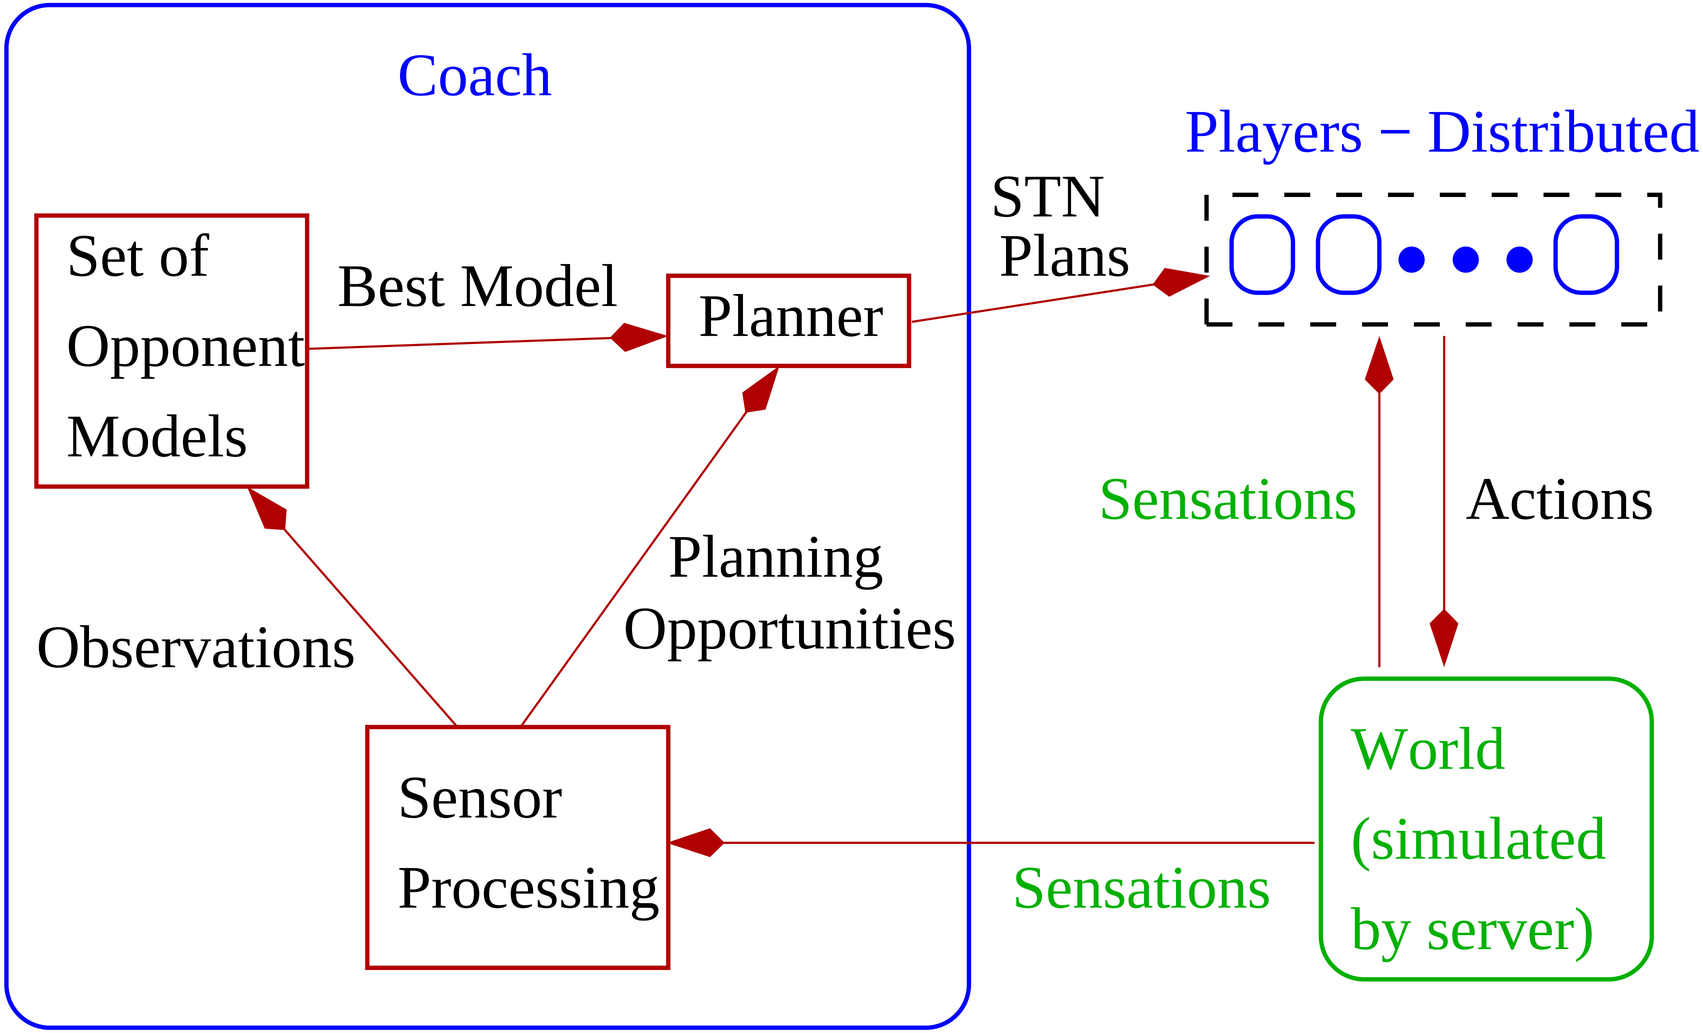

Patrick Riley's forthcoming PhD thesis

Coach-Based Arbitration

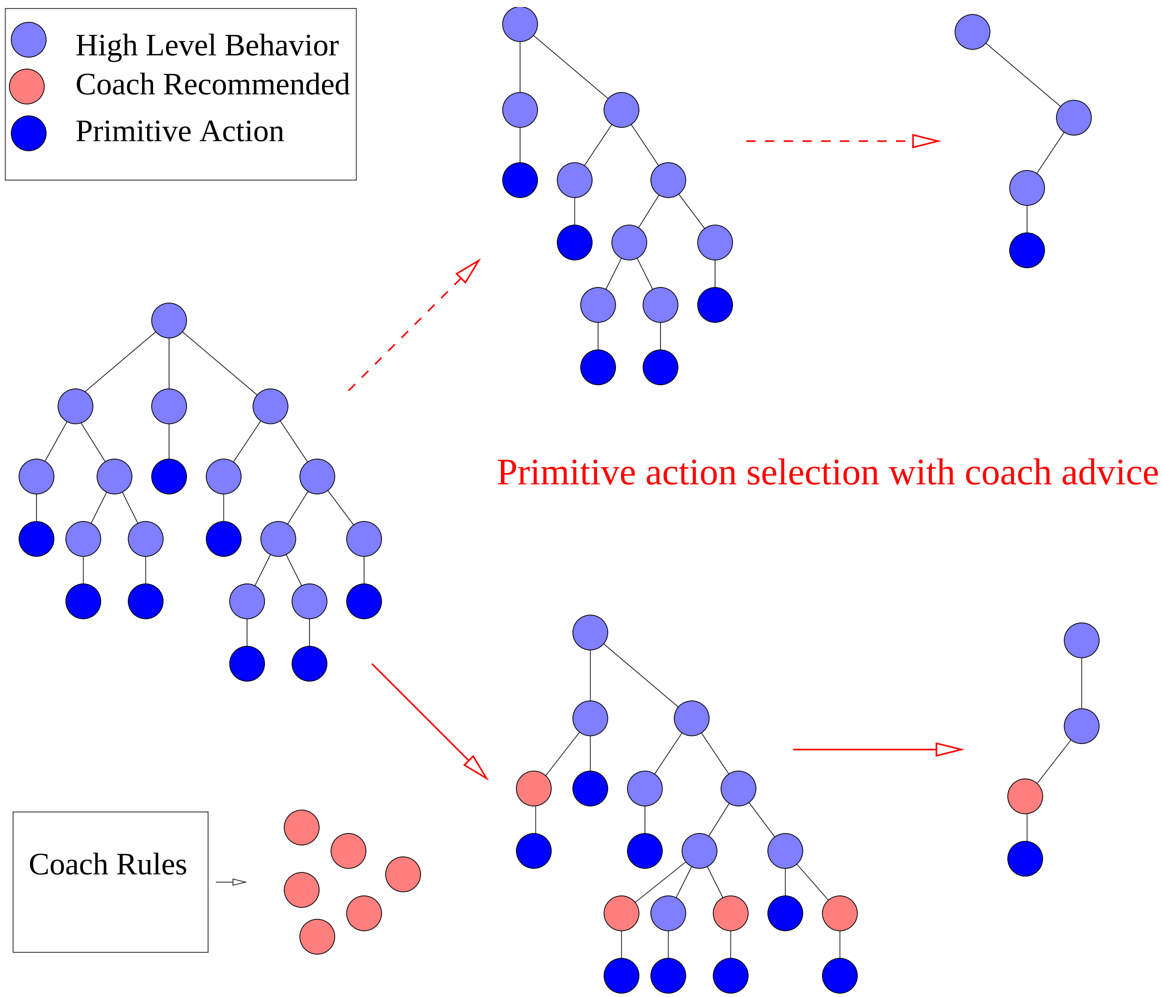

Summary: Communicating Robots

- Team perception
  - Collective building of model of the world
- Team cognition
  - Role negotiation and assignment
  - Coordination
  - Adaptation and reaction
- Team action
  - Navigation, spatial coverage

Some Relevant Publications

- Dynamic Multi-Robot Coordination, Vail, Veloso, NRL'2003.
- Communication for Multi-Robot World Model, Roth, Vail, Veloso, IROS'2003.
- Plays as Team Plans for Coordination and Adaptation, Bowling, Browning, Veloso, IJCAI'03 workshop.
- Detecting Abrupt Environment Changes, Lenser, Veloso, ICRA'2003.
- Multiagent learning with a variable learning rate, Bowling and Veloso, *Artificial Intelligence*, 2002 (IJCAI'01, ICML'01)
- Planning for distributed execution using probabilistic opponent models, Riley and Veloso, *Best Paper Award*, AIPS'02
- Integration of Advice in an Action-Selection Architecture, Carpenter, Riley, Veloso, Kaminka, RoboCup'02 Symposium
- Sensor Resetting Localization,Lenser & Veloso, ICRA'00
- Multi-fidelity Behaviors, Winner & Veloso, AAAI'00
- Efficient Robot Navigation, Bruce & Veloso, IROS'02
- Improbability Filtering, Browning, Bowling, & Veloso, ICRA'02
- Entertainment Robotics, Veloso, CACM 2002
- Recognizing probabilistic opponent movement models, Riley and Veloso, Agents'01
- Behavior Recognition, Han & Veloso, ISRR'99