Robot Dogs, Intelligent Intersections, and Connect Four: One Student's AI Research Experience

Kurt Dresner Peter Stone

Learning Agents Research Group
Department of Computer Sciences
The University of Texas at Austin

UT SURGe
December 2, 2005

Who I Am

- PhD Student in UT's CS Department
- Started in 2002

What I'm Going To Talk About

- My academic experience so far
- Projects I've worked on
- My current research
- Other stuff I'm interested in
- What I've got left to do

Why I Came To UT

- Well respected
- Liked the research I saw when I visited
- Very well-rounded program*
- Excited about Austin

First Things First

- UTCS emphasizes research from day one
- Courses (no quals)
- Picking someone to work with

Ways To Choose An Advisor

- Find someone doing something that interests you
- Find someone willing to do something that interests you

RoboCup

"By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team."

Leagues:

- Simulation
- Small Size
- Middle Size
- 4-Legged
- Humanoid (new)

RoboCup

"By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team."

Leagues:

- Simulation
- Small Size
- Middle Size
- 4-Legged
- Humanoid (new)

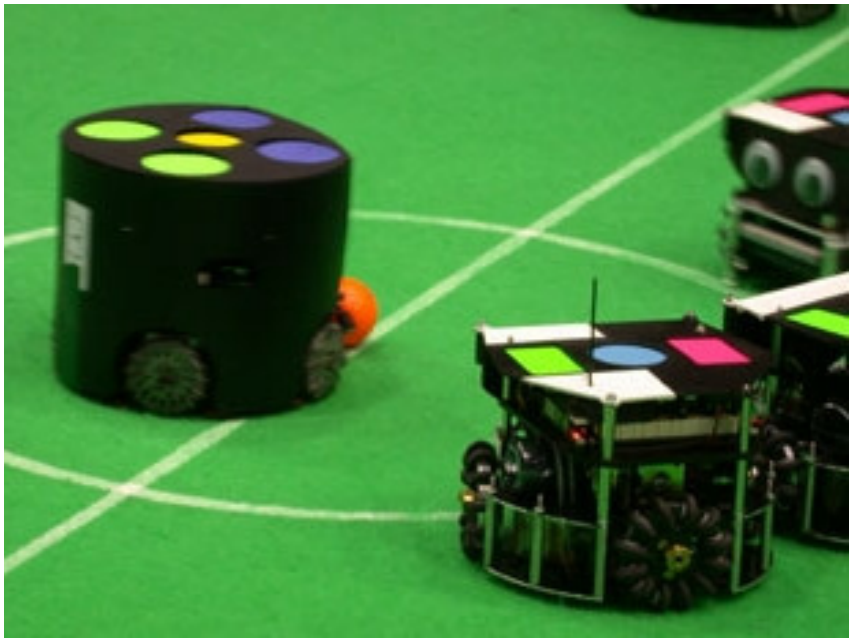

RoboCup

"By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team."

Leagues:

- Simulation
- Small Size
- Middle Size
- 4-Legged
- Humanoid (new)

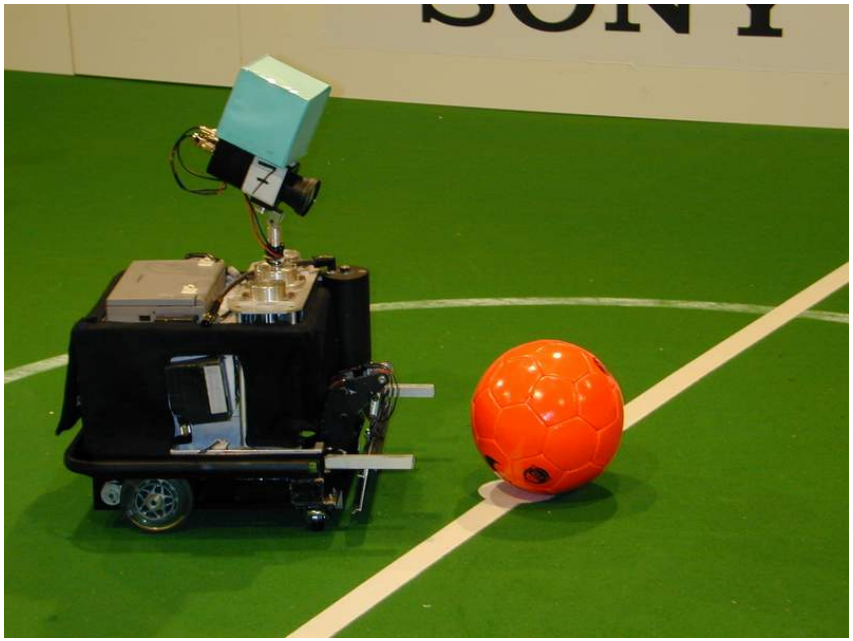

RoboCup

"By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team."

Leagues:

- Simulation
- Small Size
- Middle Size
- 4-Legged
- Humanoid (new)

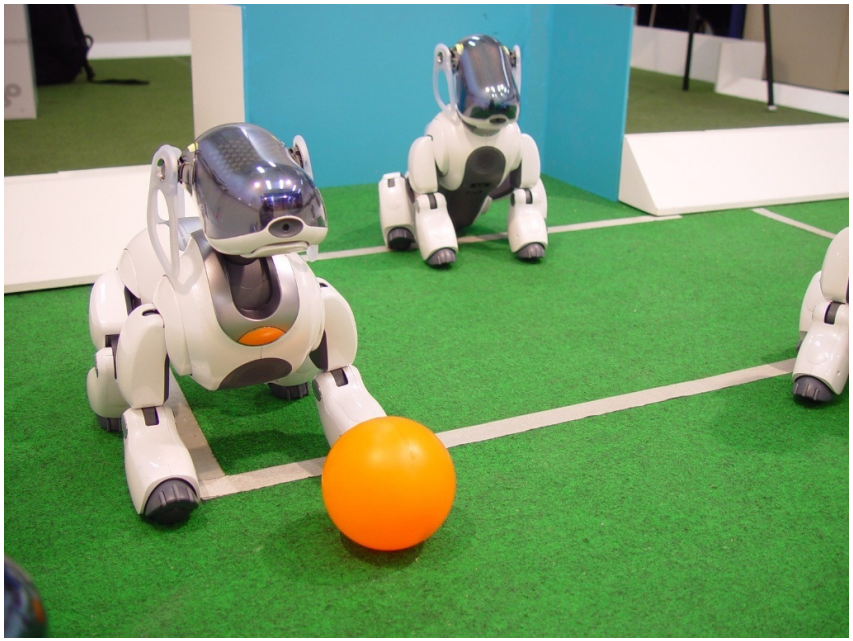

RoboCup

"By the year 2050, develop a team of fully autonomous humanoid robots that can win against the human world soccer champion team."

Leagues:

- Simulation
- Small Size
- Middle Size
- 4-Legged
- Humanoid (new)

A Member of UT Austin Villa

My Contributions

- Team coordination
- Communication

Other Parts

- Vision
- Locomotion
- Localization
- Behavior

Looking for Research

- Lost enthusiasm before finding a compelling research topic in the robotic soccer domain
- Things I found interesting about Robot Soccer
  - Autonomous agents
  - Multiagent systems (MAS)
  - Machine learning
- Entered a period of "searching" (these seem to be common in grad school)
- Eventually...

One Late Night In Austin...

- Stopped at a red light
- Driving can be fun, but usually is a chore
- What about those cars in sci-fi movies? (“Timecop”, “Minority Report”)
- What are the implications of having fully autonomous vehicles on the road?

Research Aims

Driving could be...

- Safer
  - 1.2 million deaths (WHO 1998)
  - 38.8 million injuries (WHO 1998)
- Easier
  - Age limits
  - Disabilities
  - Other impairments
- More efficient
  - Time: 46 hours
  - Fuel: 5.6 billion gallons
  - Money: $63 billion

Computers as Drivers

- More accurately sense their surroundings
- Much more precisely control a vehicle
- No distraction, road rage, drowsiness, drunkenness, aggression

Required Technology

Existing

- Intelligent cruise control
- GPS-based route planning
- Reliable wireless communication

Under Development

- Autonomous steering (Pomerleau, 1995)
- Robust lane following (Watanabe, 2005)
- Vehicle/pedestrian tracking/classification (Wender, 2005)
- High-accuracy digital maps (Weiss, 2005)

Required Technology

Existing

- Intelligent cruise control
- GPS-based route planning
- Reliable wireless communication

Under Development

- Autonomous steering (Pomerleau, 1995)
- Robust lane following (Watanabe, 2005)
- Vehicle/pedestrian tracking/classification (Wender, 2005)
- High-accuracy digital maps (Weiss, 2005)

Intersections

- Dangerous!
  - 1/3 of all accidents
  - 1/4 of all fatal accidents
- Responsible in large part for waste

Intersection Control Today

- Traffic lights

- Stop signs

- Traffic circles

A Giant Multiagent System

Automobile traffic is already a huge multiagent system.

Agents	—	Human Drivers
Mechanism	—	Traffic Signals, Stop Signs
Protocol	—	Traffic Laws

Question

New agents - should we get a new mechanism and protocol?

A Giant Multiagent System

Automobile traffic is already a huge multiagent system.

Agents	—	Human Drivers
Mechanism	—	Traffic Signals, Stop Signs
Protocol	—	Traffic Laws

Question

New agents - should we get a new mechanism and protocol?

Desiderata

- Vehicles are treated as individual agents
- Agents communicate only necessary information
- Sensor information that can be obtained with current technology
- Communication failure doesn't violate safety properties
- A simple communication protocol
- No deadlocks or starvation
- Incrementally deployable

Desiderata

- Vehicles are treated as individual agents
- Agents communicate only necessary information
- Sensor information that can be obtained with current technology
- Communication failure doesn't violate safety properties
- A simple communication protocol
- No deadlocks or starvation
- Incrementally deployable

Desiderata

- Vehicles are treated as individual agents
- Agents communicate only necessary information
- Sensor information that can be obtained with current technology
- Communication failure doesn't violate safety properties
- A simple communication protocol
- No deadlocks or starvation
- Incrementally deployable

Desiderata

- Vehicles are treated as individual agents
- Agents communicate only necessary information
- Sensor information that can be obtained with current technology
- Communication failure doesn't violate safety properties
- A simple communication protocol
- No deadlocks or starvation
- Incrementally deployable

Desiderata

- Vehicles are treated as individual agents
- Agents communicate only necessary information
- Sensor information that can be obtained with current technology
- Communication failure doesn't violate safety properties
- A simple communication protocol
- No deadlocks or starvation
- Incrementally deployable

Desiderata

- Vehicles are treated as individual agents
- Agents communicate only necessary information
- Sensor information that can be obtained with current technology
- Communication failure doesn't violate safety properties
- A simple communication protocol
- No deadlocks or starvation
- Incrementally deployable

Desiderata

- Vehicles are treated as individual agents
- Agents communicate only necessary information
- Sensor information that can be obtained with current technology
- Communication failure doesn't violate safety properties
- A simple communication protocol
- No deadlocks or starvation
- Incrementally deployable

The Reservation Idea

- Driver agents "call ahead" to reserve a region of space-time
- Intersection manager approves or denies based on an intersection control policy
- Vehicles may not enter the intersection without a reservation

Intersection Manager

- Intersection divided into an $n \times n$ grid of reservation tiles (n is the granularity)
- On a request, intersection manager simulates the journey of the vehicle through the intersection
- If no tiles occupied by the vehicle are already reserved, a reservation is granted, otherwise the request is rejected
- First come, first served

Intersection Manager

- Intersection divided into an $n \times n$ grid of reservation tiles (n is the granularity)
- On a request, intersection manager simulates the journey of the vehicle through the intersection
- If no tiles occupied by the vehicle are already reserved, a reservation is granted, otherwise the request is rejected
- First come, first served

Intersection Manager

- Intersection divided into an $n \times n$ grid of reservation tiles (n is the granularity)
- On a request, intersection manager simulates the journey of the vehicle through the intersection
- If no tiles occupied by the vehicle are already reserved, a reservation is granted, otherwise the request is rejected
- First come, first served

Intersection Manager

- Intersection divided into an $n \times n$ grid of reservation tiles (n is the granularity)
- On a request, intersection manager simulates the journey of the vehicle through the intersection
- If no tiles occupied by the vehicle are already reserved, a reservation is granted, otherwise the request is rejected
- First come, first served

A Successful Reservation

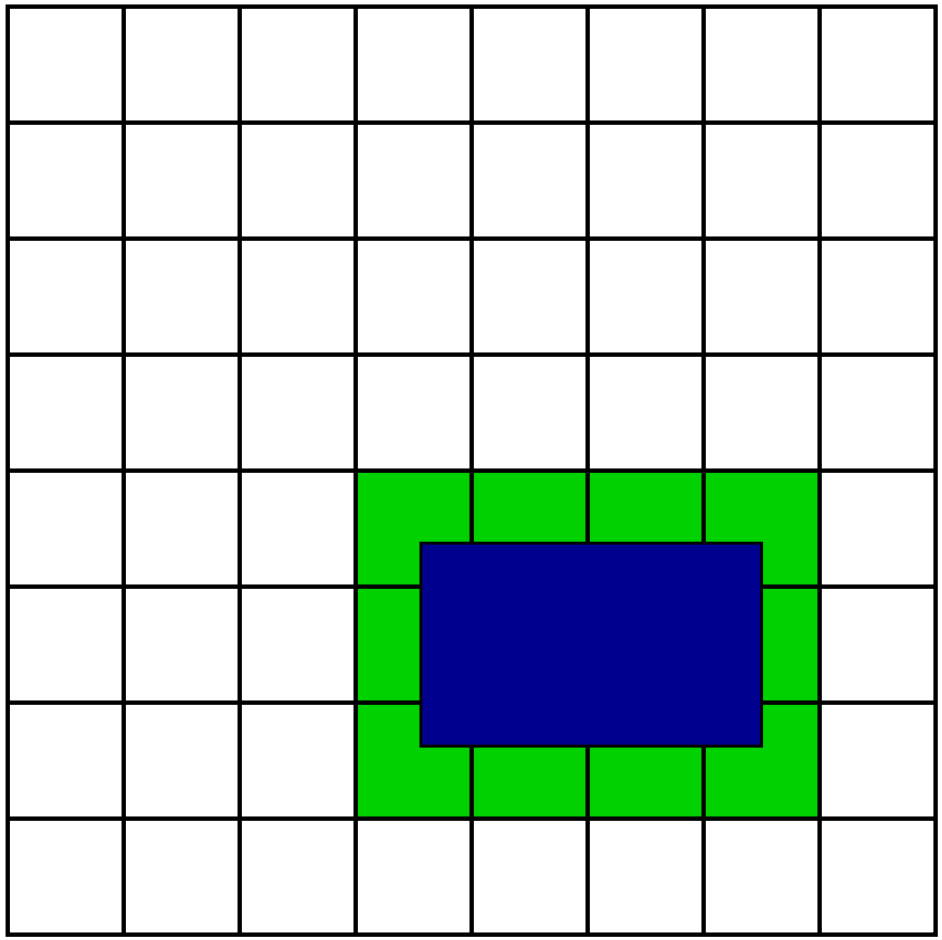

A Failed Reservation

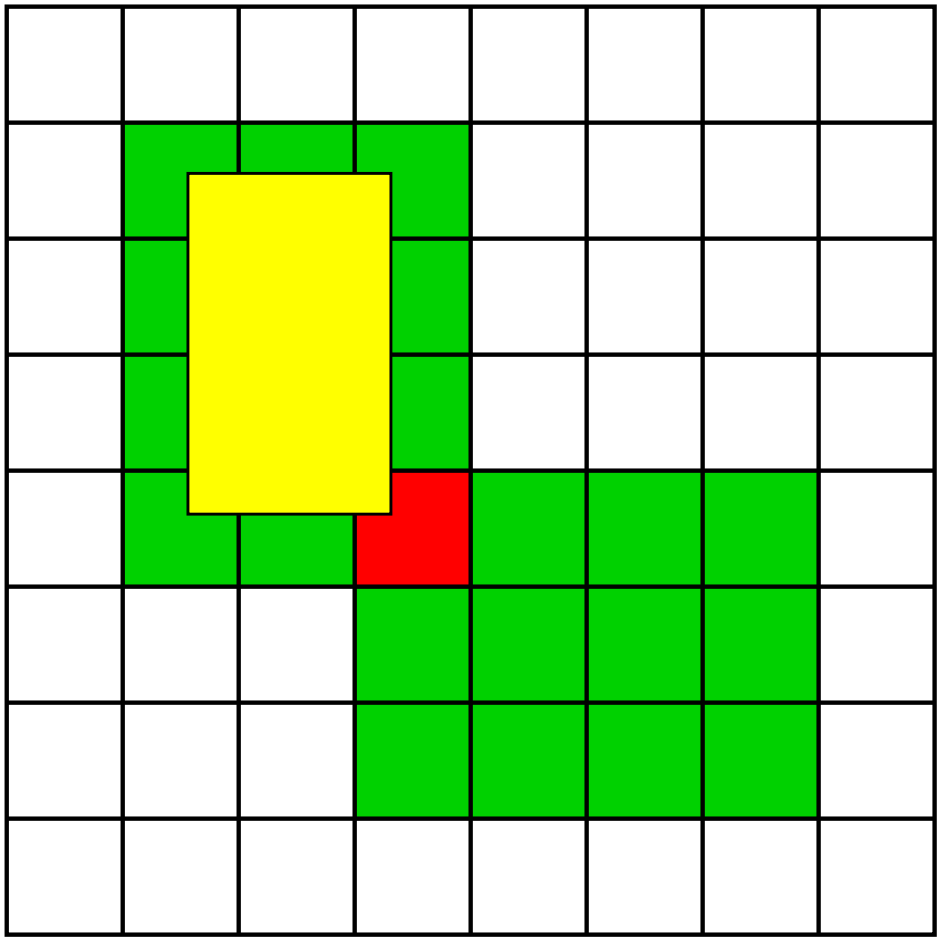

Driver Agent

- Attempt to make a reservation
- If rejected, slow down
- Try to keep reservation
- If can't keep reservation, cancel

Custom Simulator

To test our idea, we created a custom simulator.

- Simple physical model of vehicles
- Adjustable time-step (used $\frac{1}{50}$ sec)
- Amount of traffic controlled through vehicle spawning probability
- Proof of concept:
  - No turning
  - No acceleration in intersection

Implementing The Full System

Added in essential features:

- Turning
- Acceleration in the intersection
  - Disallowing causes deadlocks and starvation
  - Less efficient use of space-time
- Protocol

Implementing The Full System

Added in essential features:

- Turning
- Acceleration in the intersection
  - Disallowing causes deadlocks and starvation
  - Less efficient use of space-time
- Protocol

Implementing The Full System

Added in essential features:

- Turning
- Acceleration in the intersection
  - Disallowing causes deadlocks and starvation
  - Less efficient use of space-time
- Protocol

Protocol

- Standardized set of messages, rules
- Knowledge of internals not required
  - Different manufacturers can implement internals differently as long as they adhere to the protocol
  - New versions of software don't require updates all around
- Emulate other control policies
  - Stop sign (video)
  - Traffic light

Allowing Human Drivers and Pedestrians

- Protocol subsumes current model of traffic control
- Incorporating human drivers/pedestrians = detecting them.

A Human Approaches...

- Switch from the current reservation policy to a traffic light or other human-usable policy
- Allow the human driver or pedestrian to cross the intersection under this policy (i.e. when the light is green)
- Return to a more efficient "computers only" policy

Allowing Human Drivers and Pedestrians

- Protocol subsumes current model of traffic control
- Incorporating human drivers/pedestrians = detecting them.

A Human Approaches...

- Switch from the current reservation policy to a traffic light or other human-usable policy
- Allow the human driver or pedestrian to cross the intersection under this policy (i.e. when the light is green)
- Return to a more efficient "computers only" policy

Allowing Human Drivers and Pedestrians

- Protocol subsumes current model of traffic control
- Incorporating human drivers/pedestrians = detecting them.

A Human Approaches...

- Switch from the current reservation policy to a traffic light or other human-usable policy
- Allow the human driver or pedestrian to cross the intersection under this policy (i.e. when the light is green)
- Return to a more efficient "computers only" policy

Allowing Human Drivers and Pedestrians

- Protocol subsumes current model of traffic control
- Incorporating human drivers/pedestrians = detecting them.

A Human Approaches...

- Switch from the current reservation policy to a traffic light or other human-usable policy
- Allow the human driver or pedestrian to cross the intersection under this policy (i.e. when the light is green)
- Return to a more efficient "computers only" policy

Allowing Human Drivers and Pedestrians

- Protocol subsumes current model of traffic control
- Incorporating human drivers/pedestrians = detecting them.

A Human Approaches...

- Switch from the current reservation policy to a traffic light or other human-usable policy
- Allow the human driver or pedestrian to cross the intersection under this policy (i.e. when the light is green)
- Return to a more efficient "computers only" policy

Experimental Results

- 3 lanes
- 250m $\times$ 250m
- Speed limit: 25 m/s
- Data point: 30 min
- Granularity 24

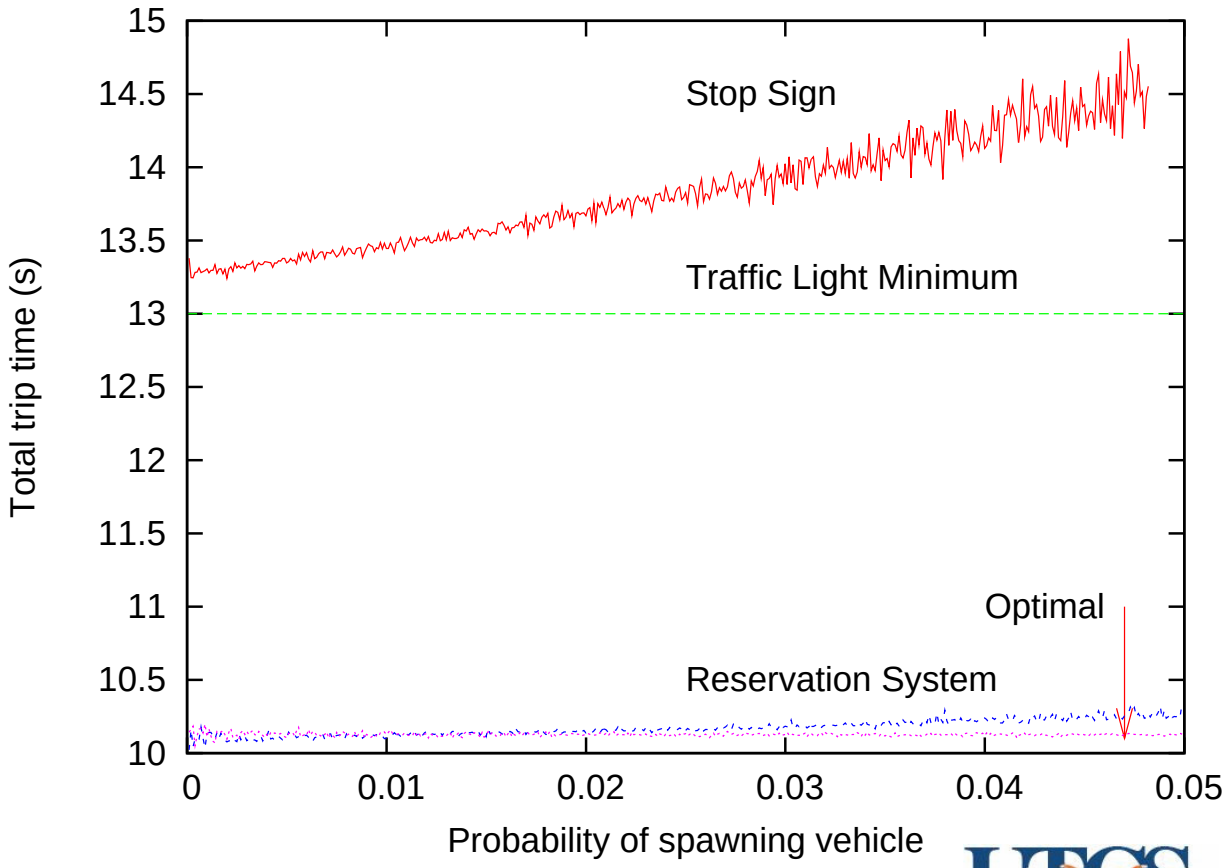

Future Work

- Lane-changing
- Multiple intersections
- Vehicle cooperation/platooning
- Multiagent learning
- Intersection as a market

Research Summary

- Intersections can be blamed for many traffic woes
- Autonomous vehicles suggest an overhaul of current mechanism for controlling intersections
- Created an intersection control mechanism using a MAS-based approach
- Outperforms traffic light, stop sign, nearly optimal

Games

- Always liked playing/analyzing games
- Another domain that incorporates my research interests
- Many programs to play games
  - Poki
  - Deep Blue
  - TD-Gammon
  - Samuel's Checkers
- Lots of research on specific games (Alberta)
- But these only play one game...

General Game Playing

- Agent receives a game description and must play the game
  - First-order logic
  - Warm-up time
  - Perfect information
- Very different games - Connect 4, Othello, Chess, Chinese Checkers, Go, Euchre, 8 Puzzle, Nim, Tic Tac Toe, Mazes
- Part of "Transfer Learning" work

Stuff We've Done

- Implemented a player and server
- Created some games
- Competed in the first annual competition (at AAAI)
- Recognize features of games
  - Board positions, coordinates
  - Pieces
  - Successor relationships
  - "Step" variables

Parts Of A Thesis

- New algorithm/idea
- Application of the idea in several domains
- Experimentation
- Writing

My Progress Towards The PhD

1. Completed courses
2. Publications in selective conferences
3. Working on a journal article
4. Need to do a thesis proposal
5. Then to write the thesis and defend it