Research Statement

Praveen Paruchuri

As multiagent systems get increasingly deployed in many critical applications in the real world, they need to account for the complexities and uncertainties associated with real-world environments that include humans and other intelligent agents. Reasoning about and dealing with intelligent entities such as humans is a very complex and challenging task since their behavior doesn't seem to be governed by simple physical laws. This behavior gets even tougher to reason about when non-cooperative settings are considered. My long-term research goal is to develop large scale multiagent systems that can plan for real world uncertainty when non-cooperative (including hostile, unfriendly or selfish) agents or environments are involved. The applications for such multiagent systems include personal assistant agents in everyday life such as offices and homes, security agents assisting police in airports and other critical infrastructure, agents assisting in disaster response scenarios, automated negotiation agents and many others. Figure 1 shows pictorially the various domains just mentioned.

Current theoretical and technological advances allow small heterogeneous multiagent teams to operate in controlled environments. The road towards building truly autonomous large-scale multiagent systems that can operate in not so friendly environments has many challenges:

- Developing expressive models and techniques to capture real-world uncertainty. While current learning or statistical techniques can automatically learn uncertainty models satisfactorily for simple patterns/behaviors, there is no technique yet (including human generated) that comes close to modeling complex behaviors such as humans.
- Improving computational efficiency of algorithms that can reason about realistic uncertainty models. While most of the current algorithms reason efficiently about specific uncertainty models, there is a need to develop efficient techniques to use many of them in tandem to reason about complex uncertainties.
- Extending current theoretical models to consider practical constraints. Many times theoretical models are far-removed from the actual hardware limitations thus limiting their applicability in real-life.

I will now describe two challenging multiagent systems that I worked on, where the agents acting in unfriendly environments face some of the challenges I just described:

Adversarial Uncertainty: Improving Security using Multiagent Systems

Security at major locations of economic, cultural or political importance is a key concern around the world, such as airports, museums, banks, government buildings and sport stadiums. Limited security resources prevent full coverage at all times, which allows adversaries

Figure 1: Domains where agents play a critical role.

to observe and exploit patterns of selective patrolling or monitoring. Therefore, randomized patrolling or monitoring is important in order to provide some level of security. On a broader level, I developed two sets of algorithms:

- When there is no model of the adversary, I provide efficient policy randomization algorithms with quality guarantees developed using the popular decision theoretic tools namely Markov Decision Problems (MDPs) and Partially Observable Markov Decision Problems (POMDPs) [3].
- When there is a partial model of adversaries, I modeled the security domain as a Bayesian Stackelberg game. I chose this specific game representation because these games can neatly capture the following observation: A main feature of patrol domain is that the patrol agent commits to a strategy first while the adversaries get to observe the agents strategy and decide on their choice of action. I then developed the fastest solver for this game named DOBSS (Decomposed Optimal Bayesian Stackelberg Solver) for the representation used [2].

The ARMOR Software: The ARMOR (Assistant for Randomized Monitoring over Routes) [4] software is a general-purpose security scheduler that is a direct outgrowth of my PhD thesis (See Figure 2 for a snapshot of the ARMOR software). ARMOR has been successfully deployed since August 2007 at the Los Angeles International Airport (LAX) to randomize checkpoints on the roadways entering the airport and canine patrol routes within the airport terminals. ARMOR casts the security problem as a Bayesian Stackelberg game and uses the DOBSS (Decomposed Optimal Bayesian Stackelberg Solver) solver to find the optimal solution. This application got extensive coverage on numerous radio reports, KNBC News, national media outlets like Newsweek, LA Times, AP News, Homeland Security

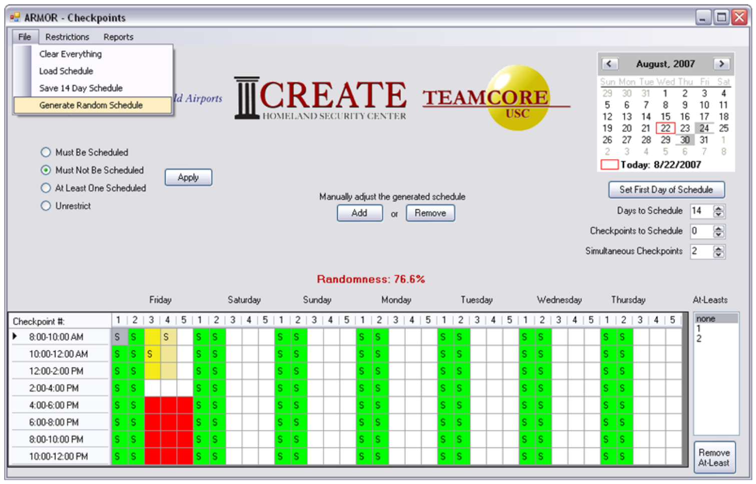

Figure 2: Snapshot of ARMOR software.

Today etc. and international newspapers like Times of India (India), Lenta.Ru (Russia) etc. ARMOR was also discussed at the Congressional committee on Homeland Security and was described as the reason for LAX being safer today. I received a commendation certificate from the City of Los Angeles thanking me for my outstanding contributions to the security of our nation. I am the first author of a book, published many papers including two finalists for best paper awards and have 2 provisional patents based on my work on improving Security. I was also officially nominated by USC in 2008 for the TR-35 award and was profiled as an innovator of this work by the USC Stevens Institute for Innovation. The success of ARMOR at LAX spawned the IRIS and GUARDS game theoretic security schedulers for the Federal Air Marshals and Transportation Security Administration respectively.

Uncertainty due to self-interested opponents: Automated negotiation

Negotiation is an interactive process by which multiple self-interested parties with limited common or uncertain knowledge about opponents try to arrive at an agreement over a set of issues with possibly conflicting preferences over the issues. The topic of negotiation has received wide attention across various fields such as political science, economics, behavioral sciences, computational sciences etc. Most work in computational modeling to-date has focused on the outcome of negotiation. My focus is to develop a general purpose negotiation agent with special emphasis on capturing the process of negotiation and not just the outcomes. To capture the evolution of the negotiation process, I modeled the negotiating agent as a dynamical system that evolves in time. In particular, the negotiation problem is modeled using a decision-theoretic framework namely Partially Observable Markov Decision Process (POMDP)[1]. This is a work in progress at CMU. I am using real human negotiation transcripts to automatically derive my POMDP agents and the aim is to make the resultant agent sensitive to human interests and biases.

Future Research

In future, I would continue to focus on the issue of building practical multiagent systems that can reason about unfriendly environments. Essentially, this involves developing better models to understand and handle uncertainty. Some of my major immediate goals include:

- Uncertainty over Adversarial Utility Models: One of the main issues I faced while building automated agents for security was how to build realistic utility functions. There are many paths that can be taken to address these issues such as using robust optimization techniques, using machine learning or other statistical techniques or simply understanding the various human factors such as motivations, cultural aspects etc.

- Uncertainty over Opponent Types: One of the domains I am currently working on is building automated negotiation agents. As mentioned earlier, this is a very rich domain for understanding the complexities involved in dealing with self-interested opponents. I am very interested in eliciting general features of a process oriented negotiation approach when dealing with a variety of opponents with uncertainty over their types.

- Robustness/Stability issues: In both the security and negotiation domains that I worked on, given that the agents work in open environments robustness and stability of the system become important issues. While I performed domain specific tests, I am interested in developing general-purpose techniques that would be valid in face of the diverse uncertainties.

References

[1] P. Paruchuri, N. Chakraborthy, R.Zivan, K. Sycara, M. Dudik, and G. Gordon. Pomdp based negotiation modeling. In *MICON Workshop, IJCAI*, 2009.

[2] P. Paruchuri, J. P. Pearce, J. Marecki, M. Tambe, F. Ordonez, and S. Kraus. Playing games for security: An efficient exact algorithm for solving bayesian stackelberg games. In *AAMAS*, 2008.

[3] P. Paruchuri, M. Tambe, F. Ordonez, and S. Kraus. Security in multiagent systems by policy randomization. In *AAMAS*, pages 273–280, 2006.

[4] J. Pita, M. Jain, J. Marecki, F. Ordonez, C. Portway, M. Tambe, C. Western, P. Paruchuri, and S. Kraus. Deployed armor protection: The application of a game theoretic model for security at the los angeles international airport. In *Industry Track, AAMAS,* **Nominee for Best Paper Award**, 2008.