Reputations and Games

Sampath Kannan

Department of Computer and Information Science

University of Pennsylvania

June 10 2010

ONR MURI Meeting

Reputation – A game-theoretic view

- Need repeated (pair-wise) interactions between agents = repeated games

- What does an agent do in a game?
  She plays a (mixed) strategy … which may change over time, depending on opponent, etc.

- Her "reputation" should be a function of this time-varying mixed strategy.

Repeated vs One-Shot Game: Example

- Prisoner's Dilemma

	Defect	Cooperate
Defect	-6	-1
Cooperate	-9	-2

Payoffs to row player;
Symmetrically to column player

- Game-theorists talk about "type" of player – defecting type or cooperating type
- Reputation: inferred type of a player based on repeated observation

To study value of reputation manager…

- Consider a 2-player game:
- Assume player Alice plays Bob repeatedly
- New twist: A does not know the payoff matrix
- Instead she must balance exploration and exploitation to minimize regret
- Exploration: A seeks to learn new matrix entries
- Exploitation: A seeks to profit from known entries
- Regret: A's lost payoff compared to the situation where she knows the matrix to start with

If A knew B's type …

- In zero-sum games she can identify her optimal strategies and learn payoffs for them with very little regret - O(n) regret where n is the number of strategies available to her. (This is best possible!)

- If instead she doesn't know B's type, best we can do currently is $O(n^2)$ regret. Probably can't be beaten in general.

- Thus in this model, a reputation manager makes a big difference!

More ways RMs can help

- Use player's behavior to estimate distribution of their types and their utilities for various outcomes
- If we assume a Bayesian prior on types, then we can design mechanisms [Z. Huang and Bei] that
  - Cause players to tell the truth about their types
  - Produce approximately optimal social welfare in some important resource allocation problems
- Reputation managers can give us these Bayesian priors

More on Zhiyi-Bei result

- Mechanism design in Bayesian model for many optimization problems:

  – Combinatorial resource allocation

  – Submodular maximization

- These problems are known to be difficult to solve even approximately in worst-case model

- But in the Bayesian model (realized by having a reputation manager) there are good approximate solutions

Future Work

- Explore other possible definitions of reputation in the game-theoretic context
  - Altruism: Each player has an altruism parameter a that corresponds to their reputation
  - Player gets payoff which is their own + a*(everybody else's)
  - Player's reputation is a and must be discovered
- In games with multiple equilibria: reputation is a way of describing which equilibrium a player prefers.

Conclusions

- Game-theoretic frameworks can be used in two stages of our Trust Management infrastructure:
  - Defining and computing reputations
  - Making decisions based on these reputations
- While problems are still challenging, intractable problems assuming a worst-case adversary, could become tractable in this game-theoretic setting
- To use this idea one needs a more complete set of models for adversaries seeking to optimize their own objective functions