

Algorithms for Graphical Games

Luis E. Ortiz

MIT CSAIL

February 1, 2005

DIMACS Workshop on Bounded Rationality

Joint work with Sham Kakade, Michael Kearns, John Langford, Michael Littman and Robert Schapire

In this talk...

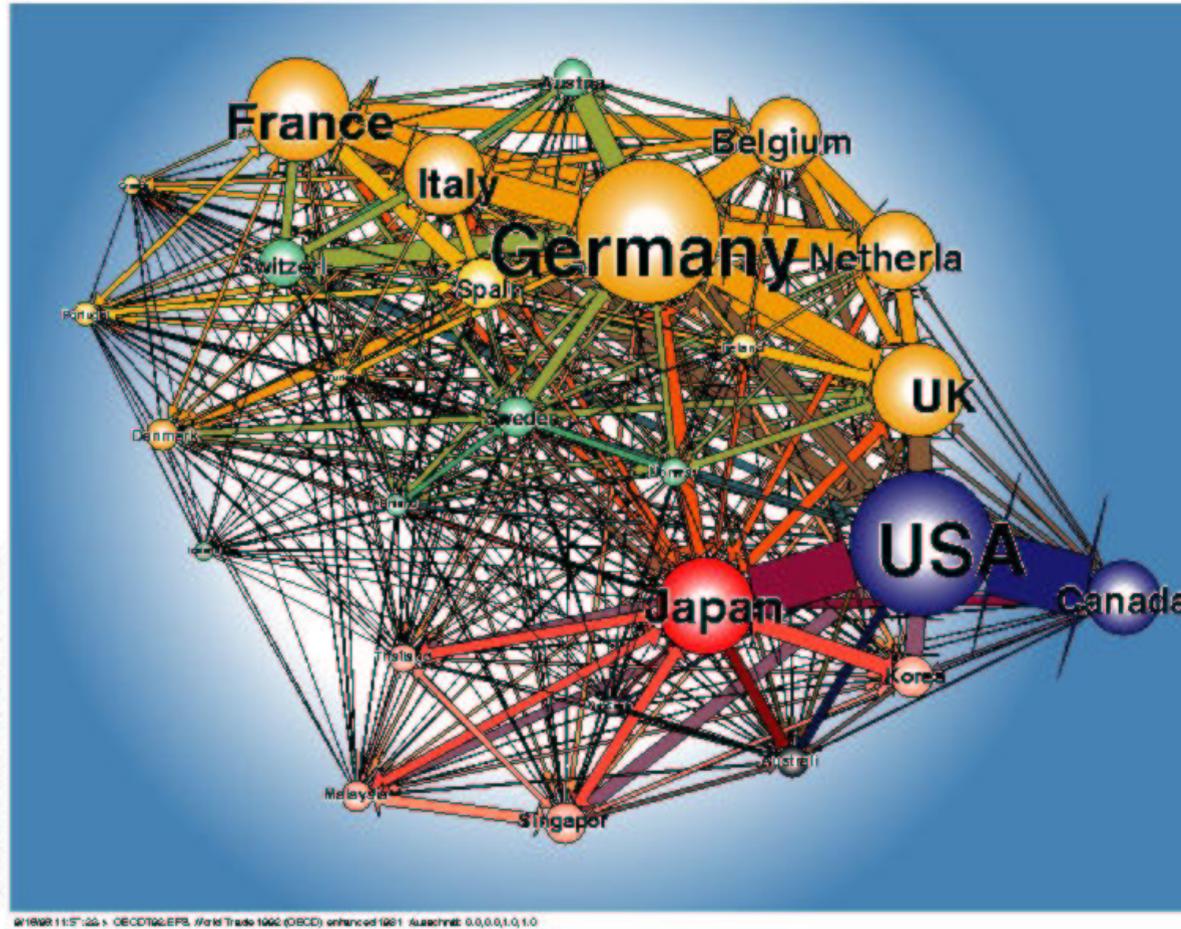
Large population games with **limited players' interaction**

Game theory

- Provides sound, rigorous mathematical formulation
- Limited attention to problem representations: commonly “flat”, large-size representations that don't exploit “structure”
- Recently, graph-based representations introduced to model interaction

Algorithms for computing equilibria in large population games with structured interactions

International Trade



(Krempel&Pleumper)

Graphical Models

Graphical models: Models of structured probabilistic interaction

- Deal with $P(X_1, \dots, X_n)$ for which naive representation is size 2^n
- Compact representations for P ; similarly for decision theory (Ex.: **Bayesian** and **Markov networks**, **influence diagrams**, . . .)
- Representation size mostly a function of the “*degree of local interaction*” among random variables X_1, \dots, X_n
- Graph allows
 - **Easy interpretation** (useful for both **modeling/knowledge-engineering** and **qualitative inference**)
 - **Compact representations**
 - **Efficient computation** in some cases

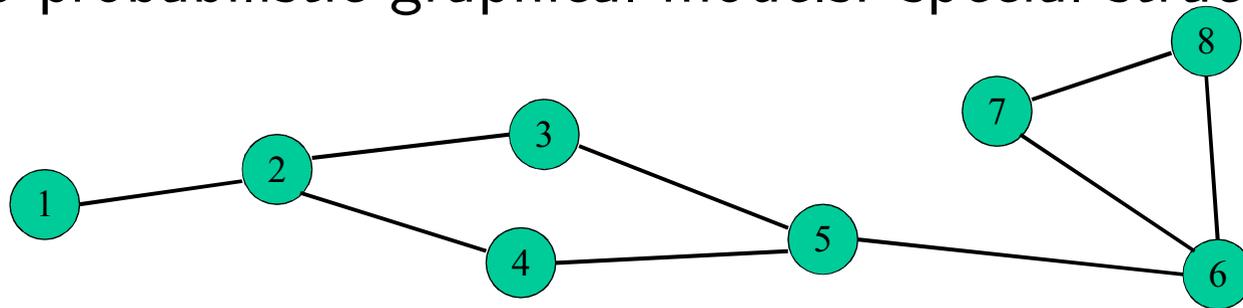
Want same benefits for game theory...

Graphical Models for Game Theory

[Kearns, Littman and Singh, 2001]

Graphical games

- Borrow representational ideas from graphical models
- Intuitive graph interpretation: A player's payoff is only a function of its neighborhood
- Ex's: geography, organizational structure, networks
- Analogy to probabilistic graphical models: special structure



Alternative graphical model formulations exist: Multi-agent influence diagrams (MAIDs) [Koller and Milch, 2001]; Game networks [LaMura, 2000]; Local-effect Games [Leyton-Brown and Tennenholtz, 2003]; Action-Graph Games [Bhat and Leyton-Brown, 2004]

What about algorithms?

- Algorithmic analogues to some inference methods in probabilistic graphical models already developed for computing Nash equilibria [Kearns, Littman and Singh, 2001; Littman, Kearns and Singh, 2001; Vickrey and Koller, 2002]
- Sometimes “tractable” computation
- Effective heuristics for general cases
- Until recently, mostly work/results on computing Nash equilibria

Overview

- **Graphical Games**
 - Compact representations for matrix/normal-form games
- Computing Nash Equilibria
- Correlated Equilibria
- Maximum Entropy Correlated Equilibria

Efficient representation and computation

Classical Example: Prisoners' Dilemma

[Tucker; see Luce and Raiffa, 1957]

- Two prisoners: To confess or not to confess???

Payoff (sentences)		Prisoner 2	
	Not-Confess	⇒	Confess
Prisoner 1	Not-Confess	1 year, 1 year	⇒ 10 years, 3 months
	Confess	3 months, 10 years	⇒ 8 years, 8 years

- Only joint **best response** is for both to **confess**!

Noncooperative Game Theory

[Von Neumann and Morgenstern, 1944; Nash, 1951]

Games with greedy players acting independently

Mathematical formulation: Normal-form games

A set of players $\{1, \dots, n\}$, each with a set of actions $\{0, 1\}$

Payoff matrix M_i : if joint-action $\vec{a} \in \{0, 1\}^n$, player i 's payoff $M_i(\vec{a})$

Mixed strategy: player i plays action 1 with probability p_i

Joint mixed strategy: product distribution

Player individually maximizes its expected payoff

Classical solution

(ϵ) -Nash equilibrium (NE) is a joint mixed strategy \vec{p}^* such that no player i can gain (more than ϵ) by unilaterally deviating from p_i^* .

NEs always exist!

Note representation size $O(n2^n)$ (exponential in number of players)

Graphical Games

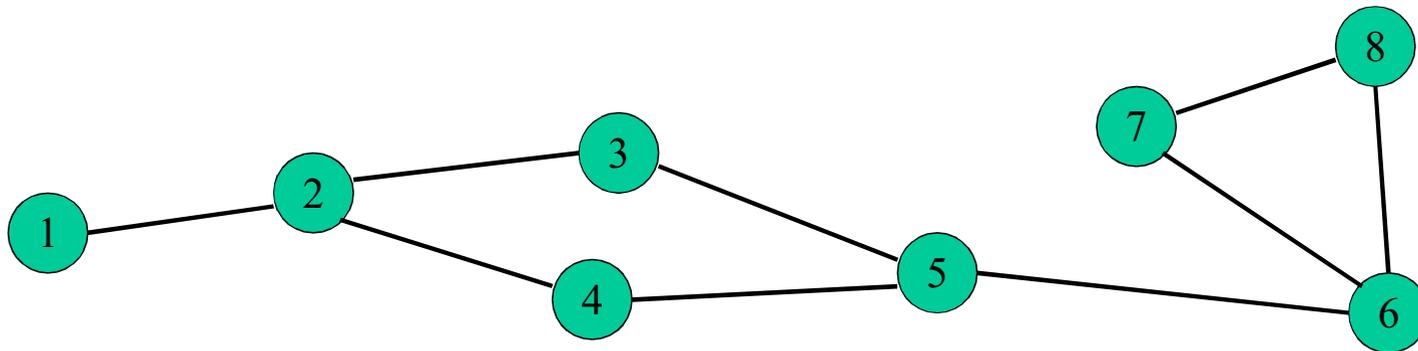
[Kearns, Littman and Singh, 2001]

Definition

- G : **undirected graph** representing the local interaction
- Player i 's payoff is only a function of its neighborhood $N(i)$
 - Implies **conditional independence payoff assumption**
- **local payoff matrix** M'_i : $M_i(\vec{a}) = M'_i(\vec{a}[N(i)])$
- **Graphical game**: $(G, \{M'_i\})$

max **degree of local interaction** $k = \max_i |N(i)| \ll n$

Representation size $O(n2^k)$ (**exponential** in max degree)



Overview

- Graphical Games
- **Computing Nash Equilibria**
 - **NashProp**: a distributed, message-passing algorithm
- Correlated Equilibria
- Maximum Entropy Correlated Equilibria

The TreeProp Algorithm

[Kearns, Littman and Singh, 2001]

Dynamic programming algorithm

Table-passing phase

$T(w, v)$ represents there exists a NE
“upstream” in which V plays v and
 W is “clamped” to w

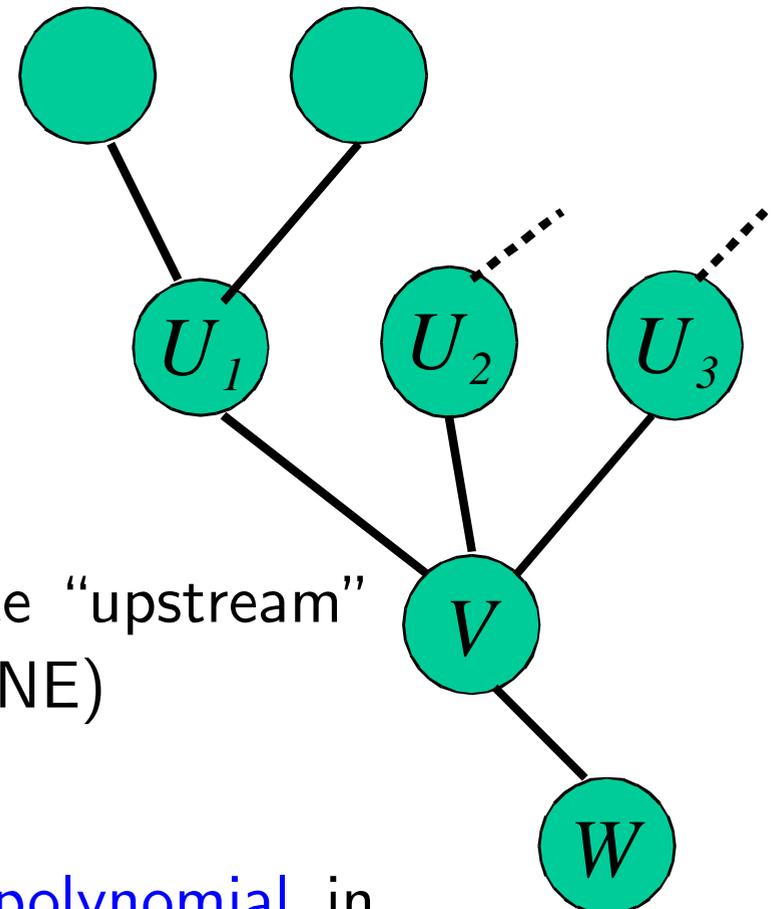
Assignment-passing phase

Assign NE mixed strategy to root.

Recursively find assignments for immediate “upstream”
neighbors consistent with tables (*i.e.*, are NE)

Representation results

- For ϵ -NE need τ -size grid for tables **polynomial** in
model size and $1/\epsilon$



The NashProp Algorithm

[Ortiz and Kearns, 2003]

A distributed, message-passing algorithm: natural extension of **TreeProp** to arbitrary graphs

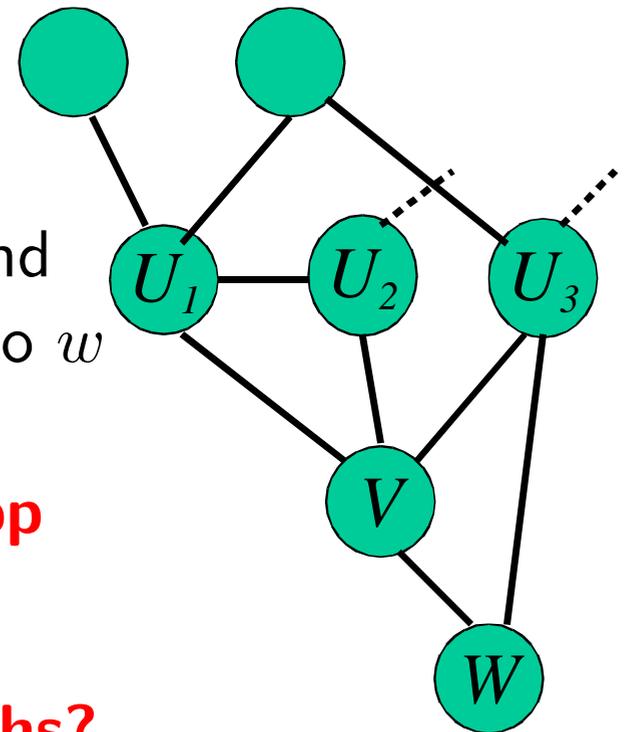
Extension as from polytree algorithm to belief propagation [Pearl, 1988]

- In our case: propagate “conditional Nash equilibria”

Table-passing phase

$T(w, v)$ represents V 's “belief” in a NE

(in the rest of the graph) in which V plays v and W is “clamped” to w



Representation results carry over from **TreeProp**

Convergence for loopy graphs?

Convergence of Table-Passing Phase

- Table-passing phase always converges
- All NE preserved
- For discretization scheme
 - Tables converge quickly (number of rounds polynomial in model size)
 - Each round takes polynomial in model size (for fixed grid size)

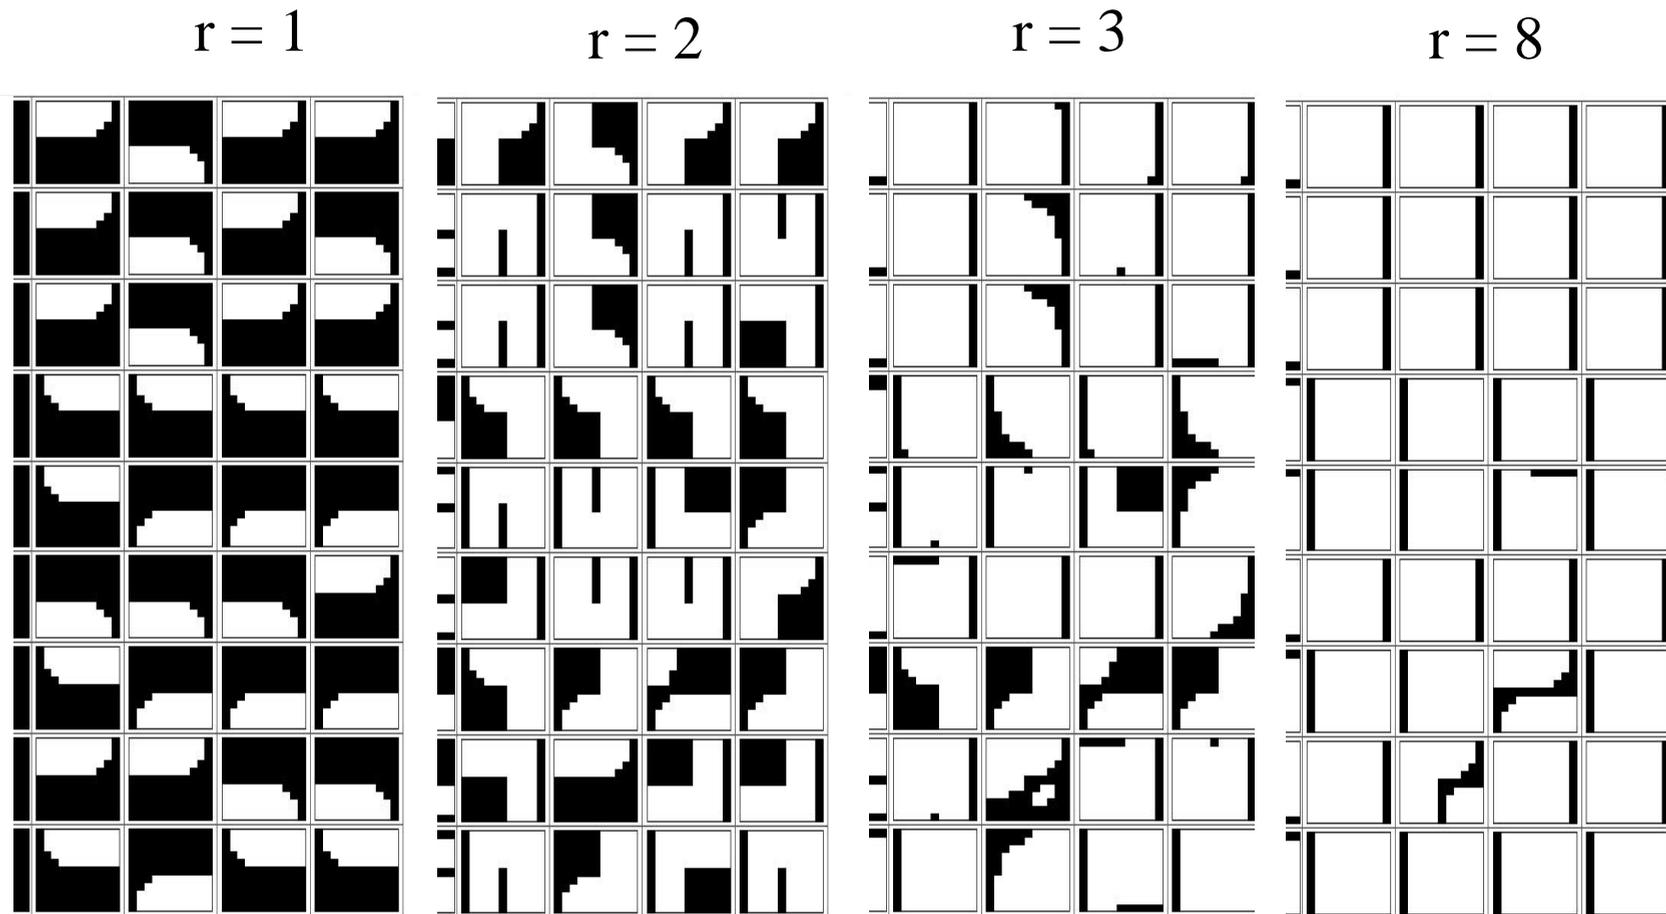
Assignment-passing phase

- NEs preserved but search still needed
- More 0's in tables can lead to significantly reduced search space
- Many heuristics possible
 - Backtracking local search
- Discretization scheme: Computation time per round **polynomial** in model size (for fixed grid size)

Discretization scheme leads to constraint satisfaction problem (CSP) formulations [Vickrey and Koller, 2002]:

NashProp is a particular instantiation of arc-consistency followed by backtracking local search in a particular CSP

Example of Ideal Behavior



- Graph: 3×3 “wrapped-around” grid
- Each row shows outbound tables for each player

NashProp

- Converging first phase with
 - table sizes for ϵ -NE polynomial in the size of the model
 - running time also polynomial (for fixed k)
- Second phase is a backtracking local search
- For both phases,

each round polynomial in the size of the model

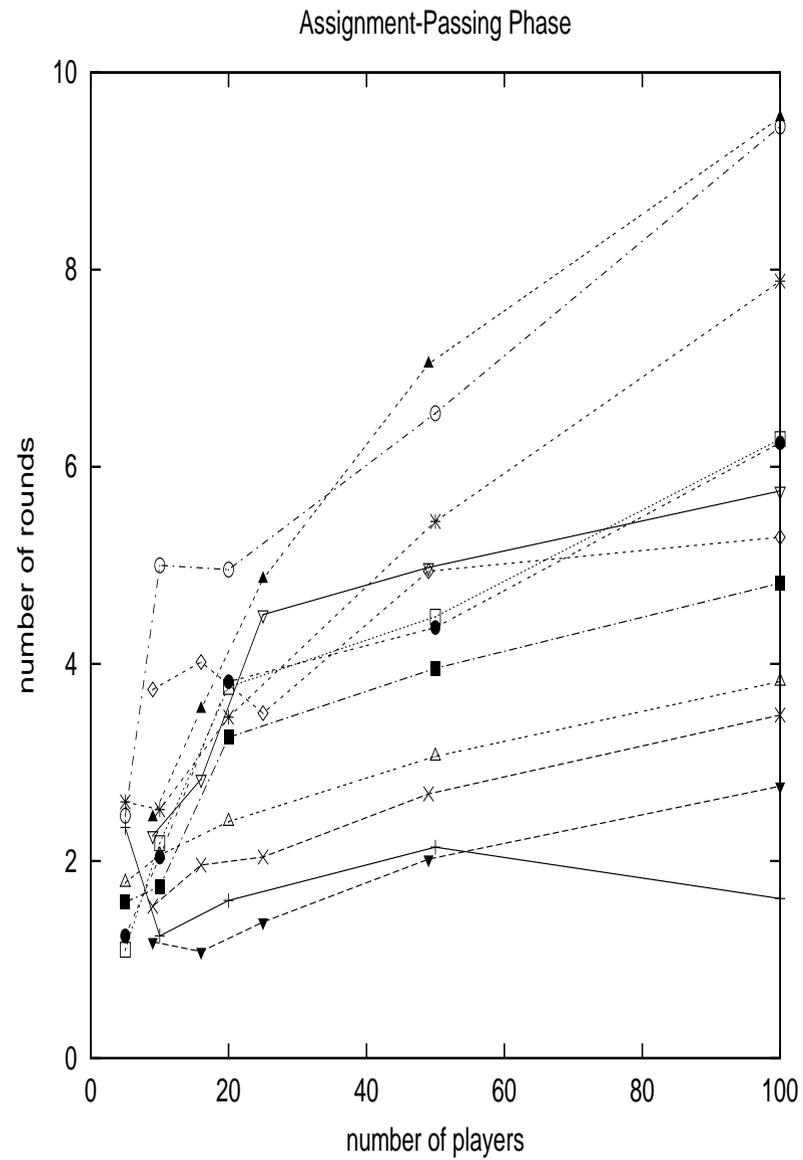
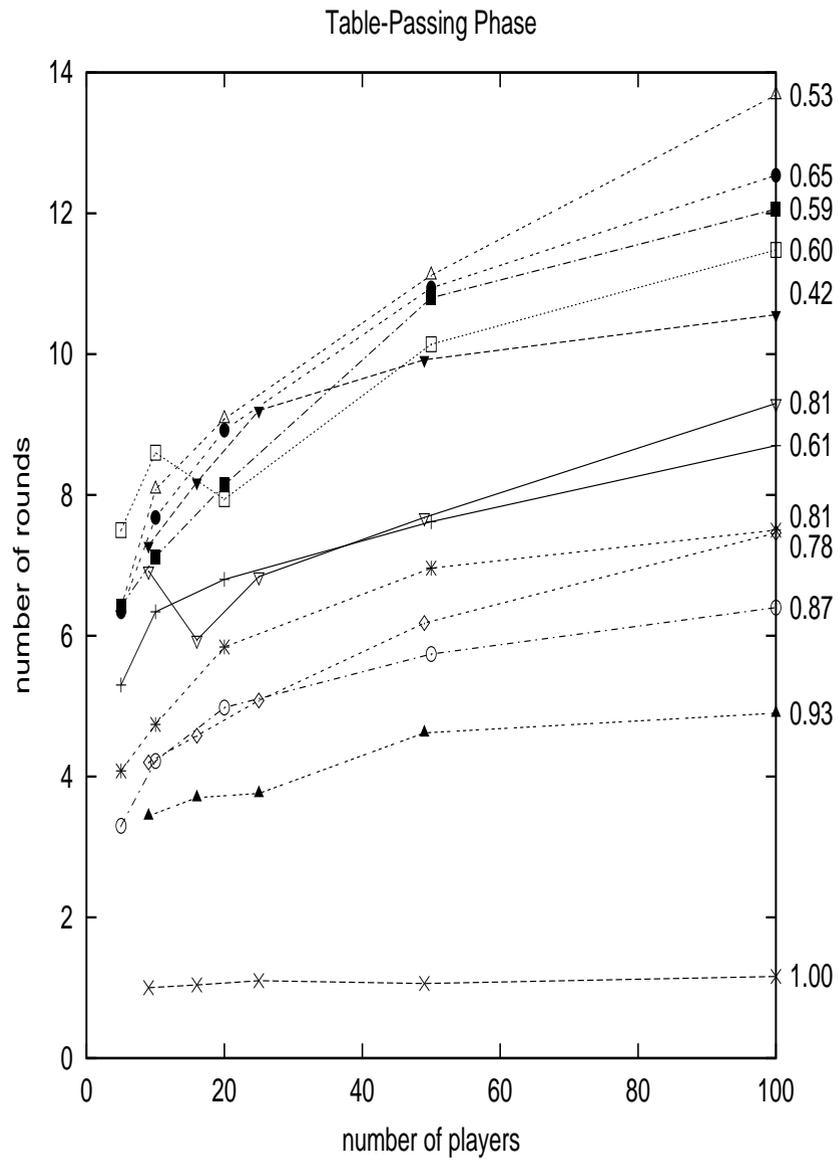
(for fixed k)

Experimental Setup

- Experiments
 - (Large) Number of players
 - Loopy graph topology
 - Random local (payoff) matrices
 - Different payoff structures

- Used heuristic local search as assignment-passing phase

Experimental Results



- cycle —+
- grid —x—
- chordal(0.25,1,2,3) —*—
- chordal(0.25,1,1,2) —□—
- chordal(0.25,1,1,1) —■—
- chordal(0.5,1,2,3) —○—
- chordal(0.5,1,1,2) —●—
- chordal(0.5,1,1,1) —△—
- grid(3) —▲—
- grid(2) —▽—
- grid(1) —▼—
- ringofrings —◇—

NE in GG: Related Work

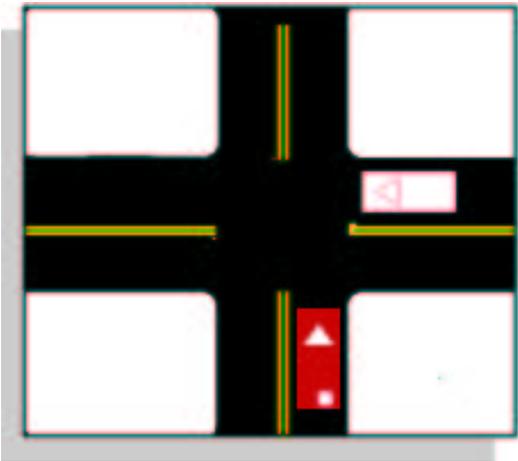
- Exact NE Computation
 - All NE in trees: exponential in representation size [Kearns, Littman and Singh, 2001]; Single NE in trees: polynomial for 2-action [Littman, Kearns and Singh, 2002], m -action open!; Single NE in loopy graphs: continuation-method heuristic [Blum, Shelton and Koller, 2003]
- Other approximation heuristics
 - CSP formulation: Cluster [Kearns, Littman and Singh, 2001] and junction-tree [Vickrey and Koller, 2002]; Gradient ascent and “hybrid” approaches [Vickrey and Koller, 2002]
- Some recent results on computing NE for torus-like GG [Daskalakis and Papadimitriou, 2004]

Overview

- Graphical Games
- Computing Nash Equilibria
- **Correlated Equilibria**
 - Definition and motivation
 - Exploiting strategic structure
 - Representation
 - Connection to probabilistic graphical models
 - Computation
- Maximum Entropy Correlated Equilibria

Example: Road Intersection

[Owen, 1995]



Two cars get to an *uncontrolled* intersection at the same time; each driver decides to stop (S) or go (G)

Payoffs for red driver (similarly for white)

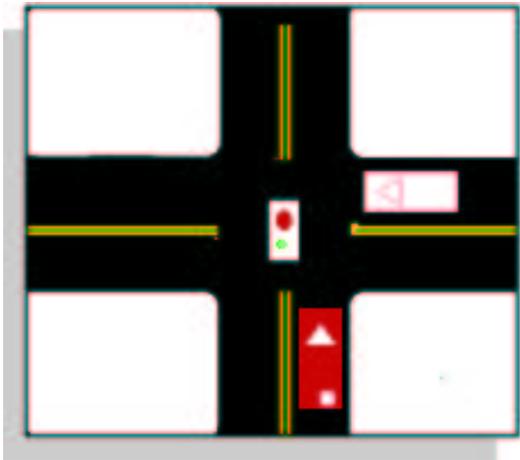
		white	
		S	G
red	S	-1	0
	G	5	-10

(NJDMV)

NE: (two pure) (1) white stops, red goes; (2) red stops, white goes; or (3) both have same “mixed strategy” over {stop, go} (both get same BUT negative expected payoff!)

Classic Example: Traffic Light

[Owen, 1995]



Traffic light acts as a **randomization device!**
Drivers respond to individual signal
Say traffic light works as follows:

Prob. signal		white	
		red	green
red	red	0	0.5
	green	0.5	0

(NJDMV)

Such mechanism (a mixture of two pure-NE) is a **correlated equilibrium**:
each driver best-response is to go if green and stop if red
“Fair” and limited “cooperation” (both achieve same AND positive
expected payoffs, not achievable by any NE alone) but still “game
theoretic”

Correlated Equilibria

[Aumann, 1974]

Same setting as before: Games with greedy players

Mathematical formulation: Normal-form games

A set of players $\{1, \dots, n\}$, each with a set of actions or pure strategies

$\mathcal{A} = \{0, 1, \dots, m\}$; joint-action $(a_1, \dots, a_n) \in \mathcal{A}^n$

Payoff matrix M_i : player i 's payoff $M_i(a_1, \dots, a_n)$

Correlated equilibrium (CE): A joint probability distribution

$P(a_1, \dots, a_n)$ such that

- Every player individually receives “suggestion” from P
- Knowing P , players are happy with “suggestion”

NE as a special case: P a product distribution; always exists!

CE in GG: Overview

- Representation

- $P(a_1, \dots, a_n)$ exponential in number of players n
- Does succinct graphical game representation \Rightarrow succinct CE representation???
- (Intuition: interaction due to game should govern correlations)
- Yes (under a reasonable equivalence class)

- Computation

- Normal-form games: CE computable via LP (with variables $P(a_1, \dots, a_n)$)
- Does succinct graphical game representation \Rightarrow efficient CE algorithm???
- Yes (for some interesting subclass of graphical games)

CE Equivalence

- Representation issue: In general, just considering arbitrary CE does not help representationally
 - Want to **preserve succinctness** of GG in CE representation as well
- Useful concepts (definitions):
 - Joint distributions P and Q over players actions are **expected payoff equivalent (EPE)** if expected payoff vectors same under P and Q
 - P and Q are **local neighborhood equivalent (LNE)** (wrt graph) if same distribution over each neighborhood (joint) actions under P and Q
- **LNE \Rightarrow EPE** (and there exist games s.t. not LNE \Rightarrow not EPE)

(Local) Markov Networks

- Compact representation of joint probability distributions
- Graph $G = (V, E)$
 - Vertices V correspond to **random variables**
 - **Potential functions** for each neighborhood of G

$$\forall i, \psi_i : \{\vec{a}^i\} \rightarrow [0, \infty)$$

- **Local Markov network (MN):** $(G, \{\psi_i\})$

$$P(\vec{a}) = \frac{1}{Z} \prod_{i=1}^n \psi_i(\vec{a}^i)$$

- Again, representation size **exponential** in **size of largest neighborhood**

Representation

[Kakade, Kearns, Langford and Ortiz, 2003]

- Efficient CE Representation Theorem

For every CE P for a graphical game with graph G , there exists a CE Q for the game with the properties

- Q is a local MN with graph G
- Q is expected payoff equivalent to P

- Proof idea/sketch: Maximum entropy (ME) distribution consistent with local neighborhood distributions of P satisfies those conditions

- Implication: Qualitative *probabilistic* properties of CEs in GGs

Computation: Normal-form Games

- **CE conditions** correspond to **linear inequalities** in the joint distribution values $P(a_1, \dots, a_n)$
- **Distribution constraints** also **linear** in $P(a_1, \dots, a_n)$
- **Known result**: We can compute a single exact CE for a game in normal-form in time polynomial in the representation size of the game ($O(nm^n)$) by using linear programming (LP)
- Can we **preserve succinctness** of GG representation in CE computation???

Computation: Graphical Games

[Kakade, Kearns, Langford and Ortiz, 2003]

- Variables: $\{P_i(\vec{a}^i)\}$ (local neighborhood marginals)
- Global CE constraints (can be a large set)
 - Best response: linear in $\{P_i(\vec{a}^i)\}$
 - Global consistency: ensure $\{P_i(\vec{a}^i)\}$ correspond to some proper global joint probability distribution $P(a_1, \dots, a_n)$
- Local CE constraints (polynomial number of linear constraints)
 - Best response
 - Local Marginal Distribution
 - Intersection Consistency
- In general, local consistency does not imply global consistency, BUT ...

Computational Result

- Local consistency sufficient for global consistency if the game graph is a tree (for example)
- Efficient Tree Algorithm

There exists an algorithm (based on LP) that finds a CE Q for tree graphical games with graph G in time polynomial in the representation size of the game

- Can be extended to bounded tree-width graphical games
- Q is also a local MN with graph G
- Can sample uniformly from set of CE in polynomial-time for bounded tree-width GG

CE in GG: Related Work

- Papadimitriou and Roughgarden [2004]
 - Alternative polynomial-time algorithm for computing a single exact CE on bounded tree-width GG
 - Also applies more generally to a larger class of “compact games”
 - Deciding whether the social-optimal CE achieves a value above some given fixed value in arbitrary GG is NP-complete
- A surprising new development [Papadimitriou, 2004]

There exists a polynomial-time algorithm for computing a single exact CE in *arbitrary* GG!

Overview

- Graphical Games
- Computing Nash Equilibria
- Correlated Equilibria
- **Maximum Entropy Correlated Equilibria**
 - Connections to learning?
 - Representation
 - Computation
 - Distributed game-theoretic interpretation

Connections to Learning?

- Natural learning dynamics converge to (the set of) CE
[Foster and Vohra, 1999; Hart and Mas-Colell, 2000; Hart and Mas-Colell, 2002]
- Are there natural learning rules directly leading to “simple” CE that exploit the strategic structure of the graphical game?
 - CE Desiderata
 - ⌋ “compactly representable”
 - ⌋ “reasonably structured”
 - Harder question than it first appears...
- There has been growing interest in trying to characterize in *any* way the behavior of natural learning rules leading to CE

Maximum Entropy Correlated Equilibria

[Ortiz, Schapire and Kakade, 2004]

- Out of all possible CE, which is more natural or likely to arise from players' dynamics?
 - This (implicit) “selection” question not particular to game theory
- **Maximum Entropy Principle** [Jaynes, 1957]
 - Useful guiding principle in many settings (e.g., statistical estimation)
 - Helps characterize equilibria in many natural systems (e.g., thermodynamics)
 - Information-theoretic view: ME distribution is simplest by being the most “uninformative”
- Not unreasonable to think ME principle is useful to CE selection question
- MECE has appealing representational and computational properties ¹

¹For technical reasons, results only valid for ϵ -CE; ignored from now on

Representation

- Given a joint-action $\vec{a} = (a_1, \dots, a_n)$, consider player i 's gain in payoff by unilaterally switching to action k : $G_{ik}(\vec{a}) = M_i(\vec{a}[i : k]) - M_i(\vec{a})$
- MECE representation

$$P(\vec{a}) = \frac{1}{Z} \prod_{i=1}^n \exp \left(- \sum_{k \neq a_i} \lambda_{ia_i k} G_{ik}(\vec{a}) \right)$$

for some $\{\lambda_{i,j,k} \geq 0\}$ (dual variables/Lagrange multipliers)

- Interpretation of λ 's: $\lambda_{ijk} = 0$ if, knowing P , player i does not want to play k when suggested j

Representation: Summary

- MECE representation properties

MECE is a local MN wrt the graph of the graphical game

- **Compact representation:** Roughly the same representation size as that for the game itself
- **Probabilistic structure:** Exploits the strategic structure of the game!

- Can sample from MECE in polynomial-time for bounded tree-width GG
 - More generally, can use a Gibbs sampler

Algorithm: Optimizing the Dual

- Let
 - P^t distribution over joint-actions at iteration t (defined by λ^t)
 - $W_{ijk}^+(t)$ = expected value (wrt P^t) of the *gains* player i achieves by switching to k instead of playing suggestion j
 - $W_{ijk}^-(t)$ = expected value (wrt P^t) of the *losses* player i suffers by switching to k instead of playing suggestion j
- **Logarithmic-gradient algorithm:** Initialize λ^0 . Iterate on t , for all players i , and pairs of actions (j, k) , $j \neq k$,

$$\lambda_{ijk}^{t+1} \leftarrow \max(0, \lambda_{ijk}^t + \delta_{ijk}^t)$$

where

$$\delta_{ijk}^t = (1/2) \ln \left(W_{ijk}^+(t) / W_{ijk}^-(t) \right)$$

Computation: Summary

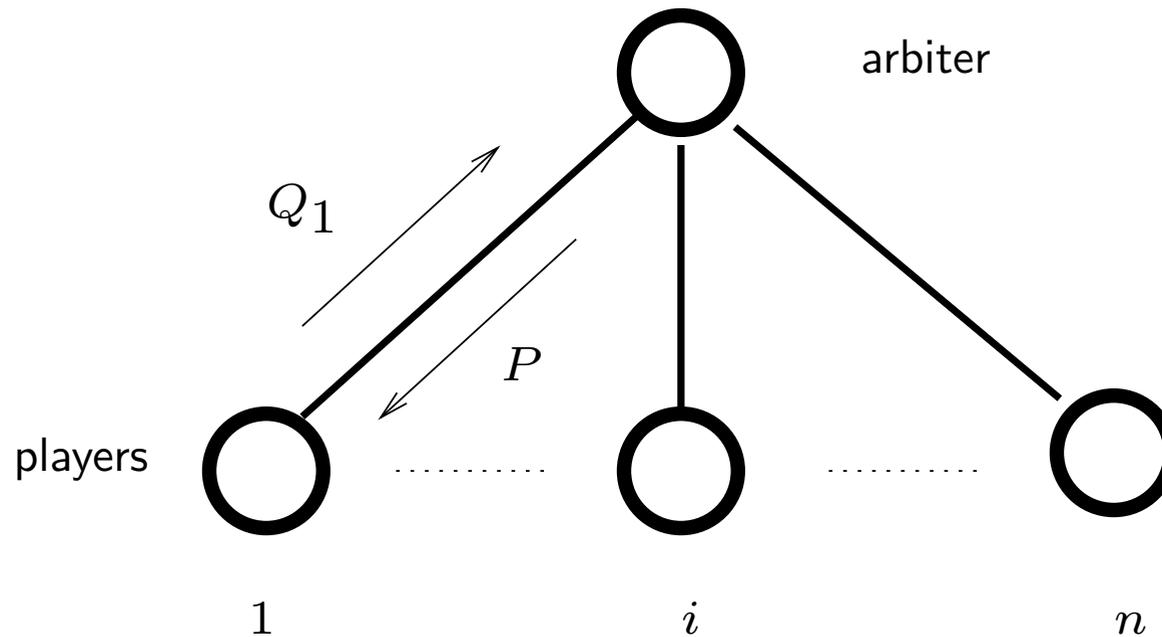
- MECE algorithmic result

Given a game \mathcal{G} , there exist gradient-based algorithms with guaranteed convergence to the MECE of \mathcal{G}

- Monotonically optimizes the dual
- Each iteration can be performed efficiently for bounded tree-width GG
 - More generally, heuristically approximate by using a Gibbs sampler
- Convergence rate still an open problem
- Strong connection to algorithms for statistical estimation: Effective in practice

Distributed Game-Theoretic Interpretation

$$P(\vec{a}) \propto Q_1(\vec{a})Q_2(\vec{a}) \cdots Q_n(\vec{a})$$



$$Q_i(\vec{a}) \propto \exp\left(-\sum_{k \neq a_i} \lambda_{ia_i k} G_{ik}(\vec{a})\right)$$

Yet, natural *learning* rules leading to MECE are still missing

NE in GG: Summary

[Ortiz and Kearns, 2002]

- **NashProp** for computing Nash equilibria in graphical games
 - Distributed, message-passing: Avoids centralized computation
 - Runs directly on game graph: Avoids operating on “hyper-graphs”
 - Generalizes approximation algorithm for trees
 - Strong theoretical guarantees for table-passing phase
 - Assignment passing phase: Simple implementations experimentally sufficient and effective in arbitrary graphs
 - Promising, effective heuristic in loopy graphs (as loopy belief propagation)

CE in GG: Summary

[Kakade, Kearns, Langford and Ortiz, 2003]

- **CE Representation Theorem**

Every CE for a graphical game can be characterized by another achieving the same expected payoffs vector and which can be represented in size polynomial in the representation size of the game

- **Efficient Tree Algorithm** (Generalizes normal-form games algorithm)

For all graphical games with graphs in a certain class, containing trees and full-graphs as “canonical” examples, we can compute a CE for the game in time polynomial in the representation size of the game

Open Problems/Areas in Graphical Games

- Computing NE for multi-action games
- Strategy-proofness of NashProp? MECE gradient-based algorithms?
- Handling parametric payoff functions: Bounded influence [Kearns and Mansour, 2002], Interdependent security games [Heal and Kunreuther, 2002; Kearns and Ortiz, 2003]
- Stochastic Graphical Games
- Learning in Graphical Games

Toolkit of representations and algorithms for large population game theory

Acknowledgements

- Dean Foster
- Stuart Geman
- Sham Kakade
- Michael Kearns
- John Langford
- Michael Littman
- Robert Schapire