CSS – TW1

International Workshop on Cooperation in Selfish Entities
(Incorporating TagWorld I)
U. Bologna, BERTINORO, Italy

COALITIONAL AND ANTAGONISTIC GAMES (ALGORITHMIC ISSUES)

Paul G. Spirakis
RACTI (Greece)

May 2006

- **ALGORITHMIC GAME THEORY**:
  A new field of Research.
- Interaction among **Selfish Entities**
- Till now, **non-cooperative strategic** games
- Point of emphasis: **Lack of coordination**
  (e.g. in routing)
- Characterization and Computing of **Nash Equilibria**.

Complex “real life” interactions require **more**:

- **Selfish coalitions** (e.g. Internet providers, politics)
- **Direct Confrontation** (e.g. security)
- **Antagonism**
  (e.g. fitness notion in biology)

•
•
•

in **static** but also in **dynamic** situations.

The algorithmic view

- **Finite** domains
- **Discrete time** dynamics
- Computability and **Efficiency**
- How to **decide**, how to **predict**
- Important parameters and concepts for measuring "**how good**"
- The computational face of **Complexity** (and how to cope with)
- Rigorous arguments

In this talk, new research on:

- Selfish Coalitions
- Direct Confrontation
- Antagonism in populations

- Emphasis in Models, analysis, proofs of statements
- Simplicity but non-triviality

The DELIS Project

EU 6th Framework Contract 001907

FET Proactive action: **Complexity**

"Dynamic and Evolving, Large Scale,
Information Systems"
(DELIS)
2005-2008

Part I

SELFISH COALITIONS

(Fotakis, Kontogiannis, Spirakis, 06)

I.I Indepentent Resources

(parallel links, machines)

m separate, identical **resources**

n **jobs** (players, users)

Each job j has an (integer) service **demand** w_j

Let $\overline{W} = \{w_1, \ldots, w_n\}$ the set of demands

Static Coalitions:

A set of $k \geq 1$ static coalitions $C_1 \ldots C_k$

is a fixed partition of $\overline{W}$ into k nonempty subsets.

I.e.

- a coalition is a **group** of jobs
- here, coalitions do not have joint members

Idea: **Each coalition is a (collective) player**

Coalitions as Strategic Games

- A **pure strategy** for a coalition C_j is the selection of a resource for each of C_j's members.
- A **mixed strategy** of a coalition C_j is any probability distribution on C_j's pure strategies.
- A **configuration** is a collection _ of pure strategies, one for each coalition
- $(\_{-j}, a_j)$ is a configuration which differs from _ only in C_j's pure strategy which is now a_j.

- **A mixed strategies profile p**
  is a collection of mixed strategies,
  one per coalition and independent
  of each other.

- The **support** of coalition C_j (in the profile p)
  is the set of pure strategies that C_j chooses
  to play with non-zero probability in p.

The payoffs

- The **Selfish Cost** of a Coalition C_j, in a configuration _, is the **maximum** demand (load) over the set of resources that C_j uses in _.

Call it $\_{j}$ (_)

E.g.

Let C_j = **3 users** with demands 7, 5, 1
Let m = **10 resources** R_1 … R_{10}
But C_j chooses in _ to put 7,1 in R_3 and 5 in R_8
Now in _, all others put 50 (total) in R_3 and 10 in R_8

So, **$\_{j}$ (_) = 58**

Now, let coalition C_j play (purely) its loads to some resources (pure strategy $\_{j}$)
Let **all other** groups play mixed strategies p.

It is reasonable to extend the **selfish cost** notion to the **conditional expectation** of the maximum demand on the resources of C_j, given that C_j has adopted $\_{j}$, and the others play p.

We call it $\_{j}$ (p, $\_{j}$)

EQUILIBRIA

- **Pure Nash Equilibrium (PNE)**

A configuration _ so that for each coalition C_j and each pure strategy a_j of C_j it holds

$$\_{j}(\_) \leq \_{j}(\_{-j}, a_j)$$

- **Mixed Equilibria**: They are mixed profiles p so that for each C_j and each $\_{j}$ in C_j's support,

$$\_{j}(p, \_{j}) \text{ is minimum}$$
$$(\text{among all } \_{j}(p, a_j))$$

- **Social Cost**: For configuration _, the Social Cost SC(_) is the maximum load over the set of all resources.
For mixed profiles p, SC(p) is the expectation of the maximum load.

Let _* a configuration that minimizes the SC(_)
We call OPT the SC(_*).

- **Price of Anarchy** : It is the maximum value R, over all NE p, of the ratio SC(p) / OPT

- **Improvement path**: It is a sequence of configurations such that any two consecutive configurations in it differ only in the pure strategy of one coalition; furthermore the cost of this coalition improves in the latter configuration.

Coalitions in Networks

- G(V,E) a directed net
- Edges have delays (nondecreasing functions of their loads)
- Each coalition is again a set of demands
- Each coalition C_j wants to route its demands from s_j to t_j
- A pure strategy of C_j is one $s_j - t_j$ path per member of C_j
- Two selfish costs:

  (i) **max** over all paths used by C_j
  (ii) **total** i.e. sum of path delays of all members of C_j

SCENARIA OF IMPROVEMENT PATHS

- The “set of resources” case
- An “improvement step” of a coalition: Its members (all) may change resources in order to reduce the coalitional cost.
- **Dynamic Coalitions**: Imagine that an arbitrary set of players (demands) forms a “temporary coalition” just to do an improvement step.
- **Unrelated resources**: The player’s demand depends also on the resource!

Pure Equilibria in the resources model.

Theorem: Even when resources are **unrelated**, even when the coalitions are **dynamic**, pure equilibria **always exist**.

Proof: Show that, starting from an arbitrary assignment, any improvement path of even dynamic coalitions of k members each has "length" (i.e. steps to reach a pure equilibrium) at most $(2k)^x / (2k – 1)$

Where x = sum over all players of the max of their demands over all resources.

We use a **generalized ordinal potential** for this.

Pure equilibria again

Theorem: Even in static coalitions (for a set of resources) and even when their number is large, it is NP – complete to find a pure Nash Equilibrium.

Small coalitions and robust equilibria

- allow **dynamic** coalitions of size k
  (say k = 2, 3 …) to be formed.
- a pure assignment is e.g. 2-robust
  if **no arbitrary coalition** of two players
  can selfishly improve its cost.

(extends to k – robust)

Note: 2-robust equilibria include all PNE of static coalitions of 2 members each.

THE ALGORITHM
"SMALLER COALITIONS FIRST" (SCF)

- **(Start)** Arbitrary assignment of users to resources
- **(loop)**

(1) Allow (in arbitrary order)
selfish improvement steps of any single player
(1-moves)

until no such move exists

(2) **If** there exists a selfish improvement "move"
of any pair of players **then** do it and go to (loop)

else we have found a 2-robust PNE

Note: Easily extends to e.g. 3-robust PNE

- Let $L_R(t)$ be the total load on resource R at "step" t.

Theorem: For identical resources the function F(t) = $\sum L_R^2(t)$ is a **weighted potential** for SCF (2)

- This assures convergence to a 2-robust PNE in at most $\frac{1}{2}$(total weight)^2 number of steps

THE PRICE OF ANARCHY

I. Coalitional chains:

Coalitions only choose from **consequtive** resources

Consider k coalitions each of $r = \frac{m}{k}$ tasks
(m = # of resources)

Assume the resources in a cycle.
Then consider the **play _**:

Each coalition chooses **uniformly** a resource
at random (as a start) and assigns its r demands
in **r consequtive resources** (e.g. clockwise)

Coalitional Chains

- _ is a Nash Equilibrium

Let M^1 = resources 1, r+1, 2r+1, …
… (k-1)r + 1 (k bins)

We have k "balls" into k bins

So,

- _ gives an anarchy ratio _ $\left(\frac{\log k}{\log \log k}\right)$
- This is, then, a lower bound to the anarchy ratio.

THE GENERAL CASE

Lower bound

Consider the play _’: (a mixed NE)

Each coalition chooses r resources uniformly at random and without replacement, and assigns one demand to each.

Lemma: For _’

(1) If k = 0 (logm / loglogm) then SC(_’)= _(k)
(2) If k = Ω (logm / log logm) then

$$SC\ (\_') = \_ \left(\frac{\log m}{\log\log m}\right)$$

Note: The random variables describing the number of coalitions hitting each resource are negatively associated.

THE GENERAL CASE

Since OPT = 1, R ≥ SC(_')

So, R ≥ _ $\left(\frac{\log m}{\log \log m}\right)$

However (see our paper)

Theorem: For every NE _,
SC (_) ≤ _ (min {k, logm/ log logm}) • OPT

So

We have a **matching** upper and lower bound for the price of anarchy. It is good when k (# coalitions) is small.

PART II

Direct Confrontation

(The price of Defense)

[Mavronicolas, Michael, Papadopoulou, Philippou, Spirakis, 06]

A strategic game on a **graph G = (V, E)**

v **attackers** each chooses one vertex to occupy

1 **defender** chooses one edge

- The payoff of an attacker is zero if it is **caught** (i.e. its vertex belongs to the defender's edge) and 1 if not caught

- The payoff of the defender is the **number of attackers** it catches.

The Price of Defense : It is the worst – case ratio (over all Nash Equilibria) of the Optimal gain of the defender (v) over the (expected) gain of the defender at a Nash Equilibrium.

Motivation: Network Edge Security [Markham, Payne, 01]
(A distributed firewall architecture)

- Are NE **tractable**?

- How does **the price of defense** vary with NE?

- How does the **structure** of G (the network) affect these questions?

- This is a confrontational game.
- **No** pure equilibria (unless the graph is trivial)
- **Def**: An **edge cover** of G is a set of edges touching all vertices of G.
- **Def:** A **vertex cover** of G is a set of vertices so that each edge of G has at least one vertex in the set.

Def: A **covering profile** is a mixed play where

(i) The Support of the defender is an Edge Cover of G

(ii) The union of the Supports of the attackers is a Vertex Cover of the subgraph of G induced by the support of the defender.

Theorem: Any NE in this game must be a Covering Profile.

Theorem: A (general) NE of this game can be found in polynomial time.

Proof

(1) Note the 1 attacker – defender game is a constant – sum game.

(2) Let $\hat{\_}$ the game with 1 attacker. Given a NE of $\hat{\_}$, let $\hat{s}$(vp) be the (sub) profile of the attacker and $\hat{s}$(ep) the (sub) profile of the defender.

Now let all attackers use $\hat{s}$(vp) independently. This is a NE of the many attackers game.

“Natural” Equilibria

(i) **Matching NE**: all attackers use a common distribution (symmetric).
All players play uniform over their support.
Each attacker uses as support an independent set of the graph.

Note: The independence number a(G) is the size of the maximum independent set of G.

Theorem: The price of defense in Matching Equilibria is a(G). They can be found in polynomial time.

(ii) **Perfect Matching NE**: If G has a perfect matching, then use this as the support of the edge player. Else, they are matching equilibria.

Theorem: Their price of defense is |V| / 2 (because any G admitting such an equilibrium must have a(G) = |V| / 2)

(iii) **Attacker symmetric NE:**

The attackers have a common support and each attacker plays uniformly on it.

Theorem: Such equilibria have a price of defense either a(G) or |V| / 2

Part III

DYNAMIC ANTAGONISM IN NETS

[Nikoletseas, Raptopoulos, Spirakis, 06]

(The survival of the weakest)

- We consider **a (fixed) network G** of **n** vertices.

- **k particles** (k ≤ n) **walk** randomly and independently around the net.

- Each particle belongs to exactly **one of two antagonistic species**, none of which can give birth to children.

- When two particles meet, they are engaged in a **“local” fight** (a small game).

- Can we **predict** (efficiently) the eventual chances of species survival?

Note:

- Classical Evolutionary Game Theory deals with multi-species competition. But its “motto” is that in each “step” any two animals are **“randomly paired”.**

This excludes any consideration of animal motions in a restricted space.

(only neighbours can interact in networks)

We examine here a simple case

- The particles are either “hawks” (H) or “Doves” (D).
- Hawks **kill** doves when they meet.
- When two hawks meet they **kill** each other
- **Doves do not harm each other** when they meet.
- **Doves are the “weakest” species**

What is their chance of eventual survival?

Note: The chance of survival of 1 dove is a lower bound to many doves.

Note: In this particular case the question is interesting when the number of hawks is even (, if we force meetings to involve only 2 particles at a time).

The "Slow" game:

- Every individual starts on a **different** vertex of G.
- At each step we choose an individual at random. **That particle then moves equiprobably to a neighbour vertex**.
- When 2 particles meet, they play the simple game

	H	D
H	(0,0)	(1,0)
D	(0,1)	(1,1)

- The process stops when only one type of individuals remains on G.

Note:

- The slow game is a **Markov process** of (at most) n^k states.

- The slow game gives the **same** survival chances as any game of concurrent moves.

- We can compute the (eventual) probabilities of the absorbing states of the process in O (n^{3k}) time. But when $k = k(n) \longrightarrow \infty$, this is too much time.

DIRECTED GRAPHS

Lemma: There are directed graphs, where the prob. of Dove survival is exponentially small.

Proof

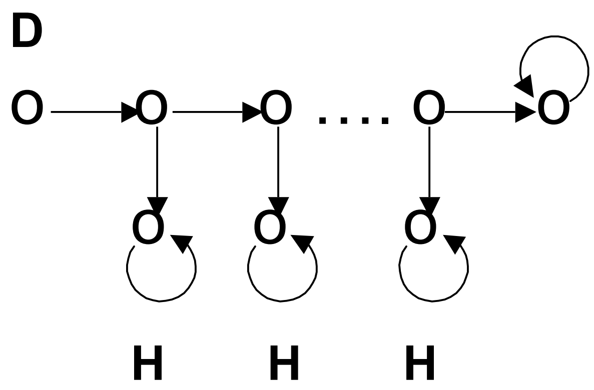

For k-1 intermediate (Chain) vertices the Dove will survive with prob $(1/2)^{k-1}$

UNDIRECTED GRAPHS

- We now concentrate on a single dove D

Def: Let **P_s(D)** be the probability of the eventual survival of the dove.

Main Theorem (undirected graphs)

Given any initial positions I of the particles and any graph (undirected) G of n vertices then

(a) We can decide in polynomial time in n if $P_s(D) = 0$ or not

(b) If $P_s(D) \neq 0$ then $P_s(D) > \frac{1}{poly(n)}$

Extinction graphs

(I) line

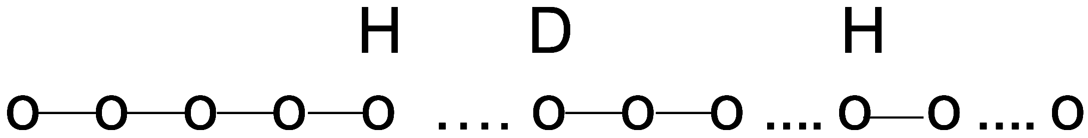

(II)

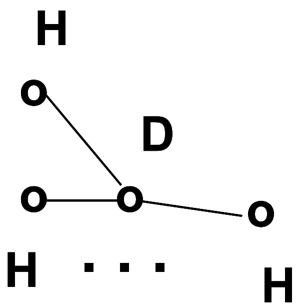

star

Lemma: the above are the only cases for which $P_s(D) = 0$

Some "easy" cases

1. **Clique** of k -1 hawks, 1 dove $P_D(s) = 1/k$

2. **Cycle** of 2 hawks 1 dove and n vertices

$$PD(s) \geq 1/n^2$$

Two “hardest” graphs

(I)

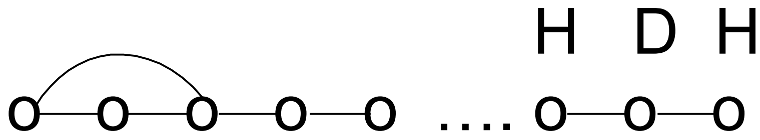

By gambler’s fortune $P_s(D) > \frac{1}{n^3}$

(II)

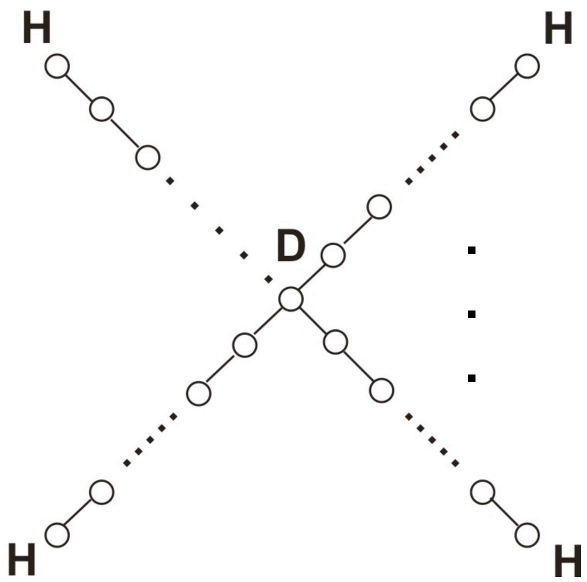

$P_s(D) > \frac{1}{n^5}$

Note: They are perturbations of the Extinction graphs.

Note: “Generalized” Gambler’s fortune

Initial treasure: The min distance from the hawks

Ruin: when it becomes zero (before the end of game)

Note: The general case reduces to the two hard graphs cases.

A simple way to estimate $P_D(s)$

- Run the game till end for N = poly (n) times
- Let x = # times the Dove survices

Then $P_D(s)$ can be estimated by $\frac{x}{N}$

Note: When N is large, the estimate becomes better.

Note: It works **only** when $P_D(s) \geq \frac{1}{poly(n)}$ else one **needs exponentially** many game simulations.

Conjecture: For directed graphs the Estimation of $P_D(s)$ is sharp P Complete.

Conjecture: Even in that case there is a polynomial time approximation scheme (reduce from **"Graph Reliability"**)

What happens when Doves reproduce?

i.e.

	H	D
H	(0,0)	(1,0)
D	(0,1)	(a,a)

and $a > 0$

Replicator dynamics:

$$\overset{0}{x}_D = x_D^2 \left(a - x_H - a x_D\right)$$

$$\overset{0}{x}_H = x_H x_D \left(1 - x_H - a x_D\right)$$

(for the **clique**)

Case 1 : a > 1
Doves dominate

Case 2 : a = 1
Anything may happen (with some probability)

Case 3 : a < 1
Hawks dominate

- **How do these extend to arbitrary graphs?**

- We examined **computational tractability questions** for coalitional and antagonistic games.

- We wish to **extend** these considerations to a **general theory**

- In all our cases, the graph was **fixed** (imposed) How about dynamic graphs?

Thank you!