One Hundred Bounds on the Price of Anarchy

Marios Mavronicolas

Department of Computer Science
University of Cyprus
Cyprus

School on Algorithmic Game Theory
Patras, December 18-19, 2006

Abstract

Algorithmic Game Theory focuses on the intersection of Computer Science and Game Theory. In the last few years, we have been witnessing a very intense research activity on this hybrid field. Computer Scientists have been working hard to compose the Theory of Algorithms and Complexity for problems originating from Game Theory.

In this talk, we will project some interesting snapshots of this composition.

Talk Structure

- Introduction
- Game Theory
- Contribution
  - The KP game
  - Discrete routing games
  - Restricted parallel links
  - Security games
  - Fair pricing games
- Research plans
- Acknowledgments

Computer Science Today

- Information Networks
- The Internet
  - Services, service providers, pricing, auctions…
  - Economic agents
  - Individual and selfish objectives
  - Competition, antagonism and lack of coordination
    ⇒main obstacle to optimization

"The Internet has arguably surpassed the von Neumann computer as the most complex computational artifact of our time."

C. H. Papadimitriou (STOC 2001)

Game Theory

- Interactions among agents
- Utilities and costs
- Rational decision-makers
- Notions of equilibria (Nash, correlated, etc.)
- Applications to Economics, Operations Research, Biology, Political Science ...

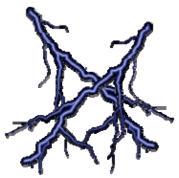

Now it is the turn of Computer Science …

Computer Science and Game Theory

- Use Game Theory to model and analyze economic phenomena over the Internet.

- Use Game Theory to analyze the selfish behavior of users in application domains (e.g. scheduling, communication, caching).

Computer Science and Game Theory (cont.)

- Import Game Theory into more traditional areas of Computer Science (e.g., Algorithms and Complexity, Networking)
  - New problems
  - New algorithmic paradigms
  - New complexity classes
- Import Computer Science into Game Theory
  - Design efficient programs for game-playing.

Computer Science and Game Theory (cont.)

- Some Concrete Tasks
  - Evaluate the performance of distributed systems with selfish entities.
  - Use game-theoretic tools to analyze specific applications:
    - scheduling
    - routing
    - P2P network creation
  - Study the algorithmic efficiency to solve computational problems in Game Theory:
    - compute Nash equilibria
    - compute Stackelberg strategies

⋮

Computer Science and Game Theory (cont.)

- Payoffs
  - A quantitative understanding of the performance of selfish distributed systems.
  - Precise models (and analytical results) for important applications.
  - A collection of
    - upper /* algorithms /*

    and

    - lower /* completeness /*

    bounds on the complexity of several algorithmic problems of Game Theory

    Most prominent: computation of Nash equilibria

Algorithmic Game Theory

- Connects Computer Science and Game Theory.
- Already two devoted conferences:
  - ACM Conference on Electronic Commerce (ACM EC)
  - International Workshop on Internet and Network Economics (WINE)
- Two forthcoming (text)books titled Algorithmic Game Theory:
  - Nisan, Roughgarden, Tardos & Vazirani (edited)

    Cambridge University Press, 2007 (expected)
  - Mavronicolas & Spirakis

    Springer-Verlag, 2007 (expected)

Game Theory Primer

- Strategic Game
  - Players
  - Strategy set for each player;
  - Utility or Individual Cost function for each player:
    - maps a strategy profile to a number
- Nash equilibrium (Nash 1950, 1951)
  - A state of the game where no player can unilaterally deviate to increase her Expected Individual Cost
- Social Cost
  - Expected Maximum Individual Cost
  - Maximum Expected Individual Cost
  - Sum of Expected Individual Costs

⋮

Price of Anarchy

$$PoA = \max_{NE} \frac{SC(NE)}{\min_{\alpha} SC(\alpha)}$$

- A prominent idea that was missing from classical Game Theory.

(Koutsoupias and Papadimitriou (STACS 1999))

Example: The KP game

Users • • •

Machines □ □ □

$1 \quad 2$

Individual Cost = Expected latency on the link she chooses

Social Cost = Expected maximum latency

$w_1 \; w_2$

(Koutsoupias and Papadimitriou (STACS 1999))

$c_1 \quad c_2$

The KP model

- Theorem 1. For the case of identical links, the fully mixed Nash equilibrium exists uniquely. For the general case, the fully mixed Nash equilibrium may only exist uniquely.

(Mavronicolas & Spirakis (STOC 2001 & Algorithmica))

The KP model

- Theorem 2. A pure Nash equilibrium always exists.

Proof Idea.

- Use lexicographic ordering on the assignments.
- Argue that the lexicographically smallest assignment is a Nash equilibrium.

(Fotakis, Kontogiannis, Koutsoupias, Mavronicolas & Spirakis (ICALP 2002))

The KP model

- Theorem 3. A pure Nash equilibrium can be computed in polynomial time.

Proof Idea.

- Use the classical Graham´s LPT scheduling algorithm for assigning weighted jobs to related machines.

(Fotakis, Kontogiannis, Koutsoupias, Mavronicolas & Spirakis (ICALP 2002))

The KP model

- Theorem 4. Computing Social Cost (=Expected Maximum) is #$\mathcal{P}$-complete.

Proof Idea.

- Use reduction from the problem of computing Pr (X$\leq$ c), where X is a sum of Bernoulli random variables.

(Fotakis, Kontogiannis, Koutsoupias, Mavronicolas & Spirakis (ICALP 2002))

CONTRIBUTION 5

The KP model

- Theorem 5. For the case of identical users and restricted to the fully mixed Nash equilibrium,

$$PoA = \Theta\left(\frac{\lg m}{\lg \lg m}\right).$$

- Theorem 6. For the case of identical links and restricted to the fully mixed Nash equilibrium,

$$PoA = \Theta\left(\frac{\lg m}{\lg \lg m}\right).$$

(Mavronicolas & Spirakis (STOC 2001 & Algorithmica))

The KP model

- Theorem 7. For the case of identical links,

$$PoA = \Theta\left(\frac{\lg m}{\lg \lg m}\right).$$

Proof Idea:

- Ball fusion:
  - it reduces the problem to the special case where all users have almost equal traffic;
  - replace two balls with their sum and assign a probability to the sum so that expected traffic for each bin is the same;
  - show that Social Cost then increases or remains the same.

(Koutsoupias, Mavronicolas & Spirakis (TOCS 2003))

CONTRIBUTION 7

The KP model

- Conjecture. The worst Nash equilibrium is the fully mixed Nash equilibrium.

(The Fully Mixed Nash Equilibrium Conjecture)

Notes:

- Intuitive and natural:
  - The fully mixed Nash equilibrium $\Rightarrow$ “collisions”;
  - Increased probability of “collisions” $\Rightarrow$ a corresponding increase to expected maximum;
- Significant:
  - It identifies the worst Nash equilibrium and trivializes the algorithmic problem of computing the worst Nash equilibrium;

(Gairing, Lücking, Mavronicolas, Monien & Spirakis (CIAC 2003 & TCS 2005))

The KP model

- Theorem 8. The Fully Mixed Nash Equilibrium Conjecture is valid for the case of arbitrary links and two identical users.

- Theorem 9. The Fully Mixed Nash Equilibrium Conjecture is valid for the case of identical users and two links.

  Proof ideas:
  - purely combinatorial
  - analytical estimations

(Lücking, Mavronicolas, Monien , Rode, Spirakis & Vrto (MFCS 2003))

The KP model

- Theorem 10. For the case of identical links, denote $h = \frac{\text{largest traffic}}{\text{average traffic}}$. Then, the Social Cost of any Nash equilibrium is no more than $2h(1+\varepsilon)$ times the Social Cost of the fully mixed Nash equilibrium.

Proof techniques:

- Uses concepts and techniques from Majorization Theory and Stochastic Orders;
- Definition of stochastic variability:
  - X is stochastically more variable than Y if for all increasing and convex functions ϕ, $\mathrm{E}(\phi(\mathrm{X})) \geq \mathrm{E}(\phi(\mathrm{Y}))$.

(Gairing, Lücking, Mavronicolas, Monien & Spirakis (CIAC 2003 & TCS 2005))

The KP model

Proof techniques (continued):

- $Y_i^j = \begin{cases} w_i, & \text{with probability } \frac{1}{m} \\ 0, & \text{with probability } 1 - \frac{1}{m} \end{cases}$

- $\widetilde{Y}_i^j = \begin{cases} \frac{W}{n}, & \text{with probability } \frac{1}{m} \\ 0, & \text{with probability } 1 - \frac{1}{m} \end{cases}$

- $\delta^j = \sum Y_i^j, \ \widetilde{\delta}^j = \sum \widetilde{Y}_i^j$

- We prove: $\max\{\delta^1, \ldots, \delta^m\}$ is stochastically more variable than $\max\{\widetilde{\delta}^1, \ldots, \widetilde{\delta}^m\}$.

- Rest uses a careful probabilistic analysis (with heavy use of Hoeffding´s Lemma).

(Gairing, Lücking, Mavronicolas, Monien & Spirakis (CIAC 2003 & TCS 2005))

The KP model

- Nashification:
  - a technique that converts a pure assignment into a Nash equilibrium with no increased Social Cost.
  - if there is a PTAS for best assignment, it can be combined with Nashification to yield a PTAS for best pure Nash equilibrium.

- Theorem 11. There is a PTAS for best pure Nash equilibrium.

(Gairing, Lücking, Mavronicolas, Monien & Spirakis (CIAC 2003 & TCS 2005))

The KP model

- Theorem 12. It is $\mathcal{NP}$-complete to approximate a worst Nash equilibrium with a ratio better than some fixed constant.

- Theorem 13. If the number of links is fixed, there is a pseudopolynomial algorithm to compute a worst Nash equilibrium.

(Gairing, Lücking, Mavronicolas, Monien & Spirakis (CIAC 2003 & TCS 2005))

CONTRIBUTION 13

Discrete Routing Games

- Unsplittable traffic
- Parallel links network
- Mixed strategies
- Convex latency functions

  $\phi(x+1) - \phi(x) \leq \phi(x) - \phi(x-1)$

- Social Cost = Sum of Expected Individual Costs

(Gairing, Lücking, Mavronicolas, Monien & Rode (ICALP 2004))

Discrete Routing Games

- Theorem 14. For discrete routing games with convex latency functions, the following hold in the case of identical users:

1. The FMNE Conjecture is valid (and the convexity assumption is essential).
2. A fully mixed Nash equilibrium may only exist uniquely.
3. There is an efficient combinatorial characterization of instances admitting a fully mixed Nash equilibrium.

(Gairing, Lücking, Mavronicolas, Monien & Rode (ICALP 2004))

CONTRIBUTION 15

Discrete Routing Games

- Theorem 15. For discrete routing games, the following are upper bounds on the Price of Anarchy:

  1. For identical users, identical links with a polynomial latency function $\phi(\lambda) = \lambda^d$, and for mixed Nash equilibria,

  $$PoA < B_{d+1}.$$

  2. For identical users, arbitrary links with polynomial latency functions and pure Nash equilibria,

  $$PoA \leq d + 1.$$

(Gairing, Lücking, Mavronicolas, Monien & Rode (ICALP 2004))

Discrete Routing Games

- Theorem 16. For identical users and arbitrary links, a pure Nash equilibrium can be computed in time $O(m \lg m \lg n)$.

The algorithm:

- It runs in $\lg n$ phases;
- In each phase, user chunks of halving size are switched together to a different machine in order to improve;
- Uses a particular data structure to implement each switch in $\Theta(\lg m)$ time.

/* proved there are $O(m)$ switches per phase /*

(Gairing, Lücking, Mavronicolas, Monien & Rode (ICALP 2004))

Discrete Routing Games

- Theorem 17. For identical users and arbitrary links with convex latency functions, an optimal assignment can be computed in time $O(m \lg m \lg n)$.

The algorithm:

- Reduction to the problem of computing a pure Nash equilibrium

/* The reduction requires convexity;

it uses a "Global Optimality = Local Optimality" –like theorem for M-convex functions. /*

(Gairing, Lücking, Mavronicolas, Monien & Rode (ICALP 2004))

CONTRIBUTION 18

Discrete Routing Games

- Theorem 18. Computing the best or the worst Nash equilibrium is $\mathcal{NP}$–complete for the case of arbitrary users, even if links are identical and their number is very small.

- Theorem 19. Counting best or worst Nash equilibria is $\#\mathcal{P}$–complete for the case of arbitrary users, even if links are identical and their number is very small.

(Gairing, Lücking, Mavronicolas, Monien & Rode (ICALP 2004))

Discrete Routing Games

- Many open problems
  - Study fully mixed Nash equilibria for the case of arbitrary users.
  - Obtain more general bounds on the Price of Anarchy.
  - Prove or disprove optimality for the algorithm to compute a pure Nash equilibrium in the case of identical users.
  - Study the approximability of best and worst Nash equilibria.

Restricted Parallel Links

- Similar to the KP model except that:
  - there is for each user a set of allowed links;
  - the cost of the user on the other links is infinite.

- Intermediate model between related links and unrelated links.

(Awerbuch, Azar, Richter & Tsur (WAOA 2003 & TCS 2006))

CONTRIBUTION 21

Restricted Parallel Links

- Theorem 20. There is a polynomial time algorithm to Nashify a given assignment for the case of restricted identical parallel links.

- Techniques and milestones:
  - it pushes the unsplittable user traffics through a flow network;
  - provides the **first** PREFLOW-PUSH like algorithm for the setting of unsplittable flows.
  - Approximation factor for optimum assignment is 2 for related links and $2 - \frac{1}{w_1}$ for identical links

(Gairing, Lücking, Mavronicolas & Monien (STOC 2004))

Security Games

- A network;
- Many attackers, one defender

⇓ choose vertices to destroy

⇓ chooses an edge to protect

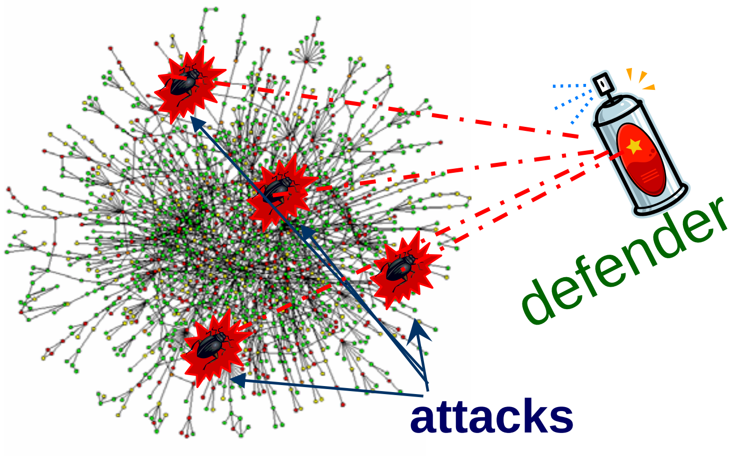

- Attacker wins the probability of not being caught
- Defender wins the expected number of attackers it catches

What are the Nash equilibria?

How can we evaluate them?

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

CONTRIBUTION 23

A Graph-Theoretic Security Game

- Associated with G(V, E), is a strategic game:

$$\Pi(G) = \langle \mathcal{N}, \{S_i\}_{i \in \mathcal{N}}, \{\mathsf{IP}\}_{i \in \mathcal{N}} \rangle$$

  - $\mathcal{N} = \mathcal{N}_{vp} \cup \mathcal{N}_{ep}$
  - ν attackers (set $\mathcal{N}_{vp}$) or vertex players vp_i
    - strategy set : $Svp_i = V$
  - a defender or the edge player ep
    - strategy set : $Sep = E$

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Individual Profits

- In a pure profile $\vec{s} = \langle s_1, \ldots, s_\nu, s_{ep} \rangle \in \mathcal{S}$
  - Vertex player Individual Profit:
    - $\mathsf{IP}_i(\vec{s}) = 0$ if $s_i \in s_{ep}$ or 1 otherwise

    1 if it selected node is not incident to the edge selected by the edge player, and 0 otherwise
  - Edge player´s Individual Profit:
    - $\mathsf{IP}_{ep}(\vec{s}) = |\{i : s_i \in s_{ep}\}|$

    the number of attackers placed on the endpoints of its selected edge

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Example

- a graph G
- ν=4 vertex players ☒
- edge player ep ☑

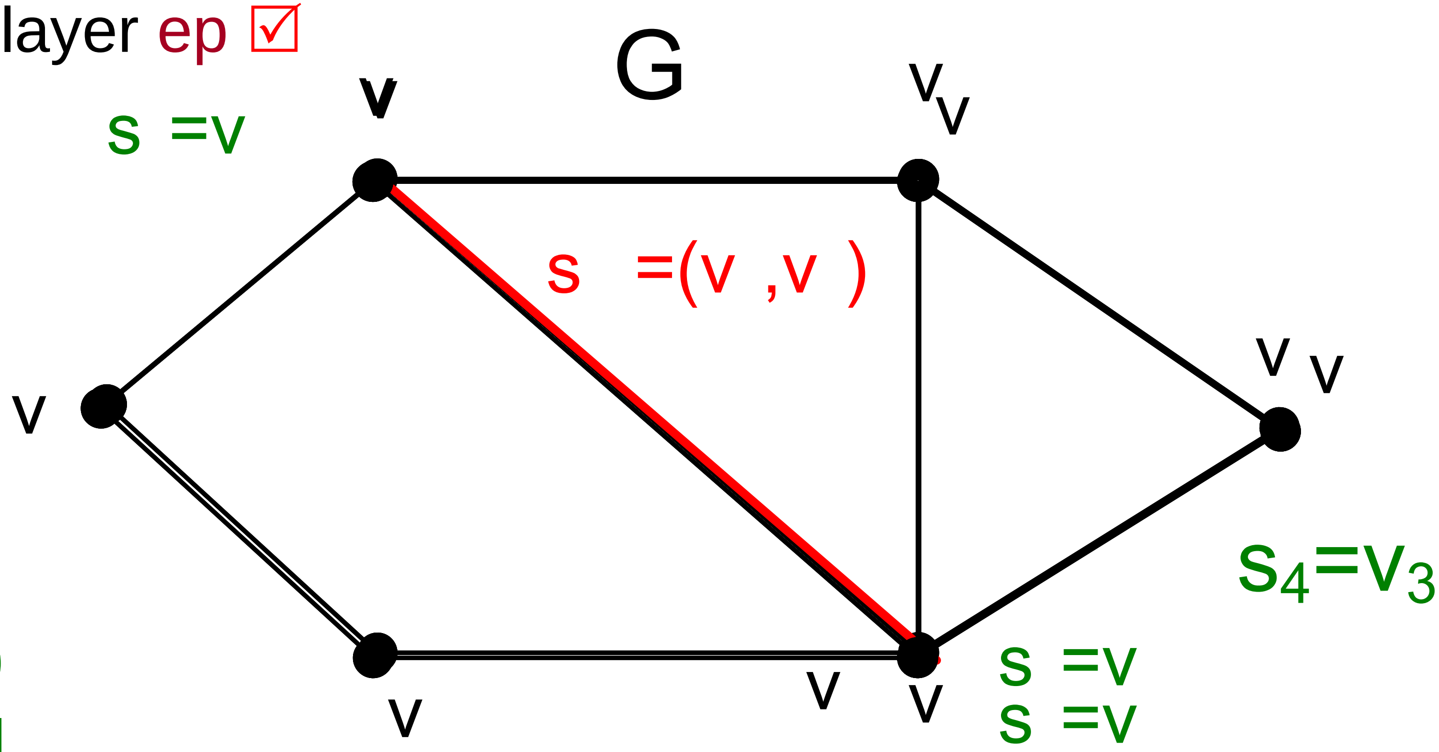

- $IP_s(ep)=3$
- $IP_s(vp_1)=0$
- $IP_s(vp_4)=1$

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

CONTRIBUTION 26

The Price of Defense

In a profile s, the Defense Ratio is:

- $DR_s = \dfrac{\text{Number of attackers}}{\text{Expected number of attackers caught in } s}$

The worst-case value of the Defense Ratio is the Price of Defense:

- $PoD_G = \max \dfrac{\text{Number of attackers}}{\text{Expected number of attackers caught}}$

- It is the Price of Anarchy applied to the security game.

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

Notation

In a profile s,

- Support of player i $\text{Support}_s(i)$
  - set of pure strategies that it assigns positive probability
- $\text{Support}_s(vp)$ = the supports of all vertex players
- $P_s(\text{Hit}(\upsilon))$ = Probability the edge player chooses an edge incident to vertex υ
- $VP_s(\upsilon)$ = expected number of vertex players choosing vertex υ
- $VP_s(e) = VPs(\upsilon) + VPs(u)$, for an edge $e=(u, \upsilon)$

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Notation (cont.)

- Uniform profile:
  - if each player uses a uniform probability distribution on its support. I.e., for each player i,

$$s_i(x) = \frac{1}{|\mathsf{Support}_{\mathbf{s}}(i)|}, \text{ for any } x \in \mathsf{Support}_{\mathbf{s}}(i).$$

- Attacker Symmetric profile:
  - All vertex players use the same probability distribution

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 29

Computational Complexity Tools (1/2)

- UNDIRECTED PARTITION INTO HAMILTONIAN CIRCUITS OF SIZE AT LEAST 6

  Input: An undirected graph G(V,E)

  Question: Can the vertex set V be partitioned into disjoint sets V_1, Λ, V_k, such that each $|V_i| \geq 6$ and $G(V_i)$ is Hamiltonian?

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 30

Computational Complexity Tools (2/2)

- We provide the first published proof that:
- Theorem 21. UNDIRECTED PARTITION INTO HAMILTONIAN CIRCUITS OF SIZE AT LEAST 6 is $\mathcal{NP}$-complete.

Proof.

Reduce from:

- the directed version of the problem for circuits of size at least 3
  - known to be $\mathcal{NP}$ -complete [GJ79] ■

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 31

Graph-Theoretic Tools (1/2)

- KÖNIG-EGENVÁRY MAX INDEPENDENT SET

Instance: A graph G(V, E).

Output: A Maximum Independent Set of G is König-Egenváry ($\alpha(G) = \beta'(G)$) or No otherwise.

- Previous Results for König-Egenváry graphs
  - (Polynomial time) characterizations [Deming 79, Sterboul 79, Korach et. al, 06]

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 32

Graph-Theoretic Results (2/2)

- Proposition 1. KÖNIG-EGENVÁRY MAX INDEPENDENT SET can be solved in polynomial time.

Proof.

- Compute a Min Edge Cover EC of G
- From EC construct a 2SAT instance ϕ such that
  - G has an Independent Set of size $|EC|(=\beta'(G))$ (so, $\alpha(G) = \beta'(G)$) if and only if ϕ is satisfiable.

■

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 33

Existence of Pure Nash Equilibria

- Theorem 22. There is no pure Nash equilibrium.

  Proof Sketch.

  Let $e=(u,v)$ the edge selected by the edge player in s.

  - |E| > 1 $\Rightarrow$ there exists an edge $(u', v')=e' \neq e$, such that $u \neq u'$.
  - If there is a vp_i located on e,
    - vp_i will prefer to switch to u and gain more

      $\Rightarrow$ Not a Nash equilibrium.
  - Otherwise, no vertex player is located on e.
    - Thus, $IC_{ep}(s)=0$,
    - ep can gain more by selecting any edge containing at least one vertex player.

      $\Rightarrow$ Not a Nash equilibrium. ■

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Characterization of Mixed Nash Equilibria

- Theorem 23. A profile s is a Nash equilibrium if and only if the following two conditions hold:
  1. For any vertex $v \in$ Support$_{\mathbf{s}}(vp)$,

     $$P_{\mathbf{s}}(\text{Hit}(v)) = \min_{v' \in V} P_{\mathbf{s}}(\text{Hit}(v')).$$
  2. For any edge $e \in$ Support$_{\mathbf{s}}(ep)$,

     $$VP_{\mathbf{s}}(e) = \max_{e' \in E} VP_{\mathbf{s}}(e').$$

- Note.

  Does not imply a polynomial time algorithm for computing a Nash equilibrium..

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Computation of General Nash equilibria

- Theorem 24. A mixed Nash equilibrium can be computed in polynomial time.

  Proof idea:

  - Reduction to a two-person, constant-sum game:
    - Consider a two players variation of the game $\Pi(G)$:
      - 1 attacker, 1 defender
    - Show that it is a constant-sum game
    - Compute a Nash equilibrium s´ on the two players game (in polynomial time)
    - Construct from s´ a profile s for the many players game:
      - which is Attacker Symmetric
      - show that it is a Nash equilibrium ■

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Necessary Conditions for Nash Equilibria (1/2)

- Proposition 1. For a Nash Equilibrium s, $\text{Support}_s(\text{ep})$ is an Edge Cover of G.

Proof Sketch.

- Assume in contrary that there exists a vertex $v \in V$ such that

$v \notin \text{Vertices}(\text{Support}_s(\text{ep}))$,

$\Rightarrow$ $\text{Edges}_s(v) = \varnothing$ and $P_s(\text{Hit}(v)) = 0$.

$\Rightarrow$ any vertex player vp_i chooses some such v with probability 1,

$\Rightarrow$ $VP_s(e) = 0$,

$\Rightarrow$ $IP_s(\text{ep}) = 0$.

Since s is a Nash equilibrium, $IP_s(\text{ep})>0$.

$\Rightarrow$ A contradiction. ■

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

CONTRIBUTION 37

Necessary Conditions for Nash Equilibria (2/2)

- Proposition 2. For a Nash Equilibrium s, $\text{Support}_s(vp)$ is a Vertex Cover of the graph $G(\text{Support}_s(ep))$.

Proof Sketch.

- Similar to Proposition 1.

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Covering Profiles

- Definition. A Covering profile is a profile s such that
  - $\text{Support}_s(ep)$ is an Edge Cover of G
  - $\text{Support}_s(vp)$ is a Vertex Cover of the graph $G(\text{Support}_s(ep))$.

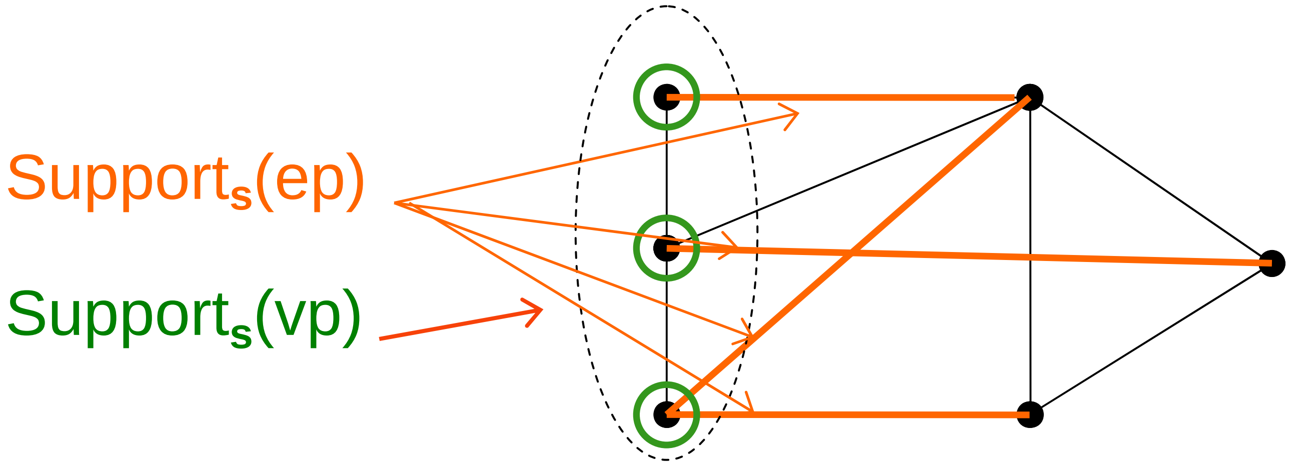

- Proposition 3. A Nash equilibrium is a Covering profile.

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Independent Covering Profiles

- Definition. An Independent Covering profile **s** is a uniform, Attacker Symmetric Covering profile s such that:
  1. $Support_s(vp)$ is an Independent Set of G.
  2. Each vertex in $Support_s(vp)$ is incident to exactly one edge in $Support_s(ep)$.

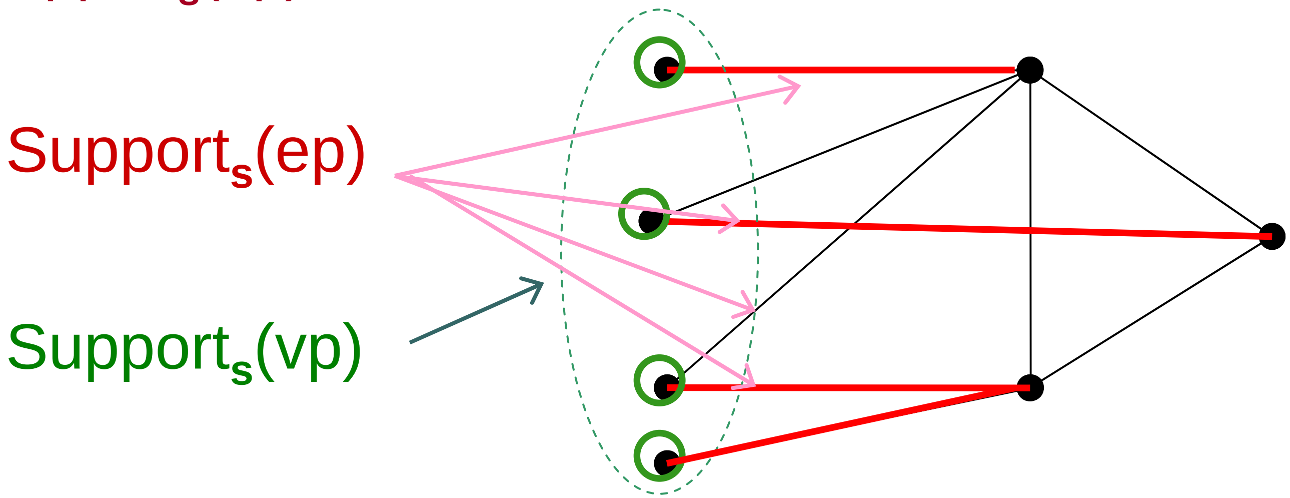

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Matching Nash Equilibria

- Proposition 4. An Independent Covering profile is a Nash equilibrium, called Matching Nash equilibrium.

Proof.

We prove Conditions (1) and (2) of the characterization of a Nash equilibrium:

- Condition (1): Consider a vertex $v \in$ Support$_s$(vp).
  - v is incident to exactly one edge $e \in$ Support$_s$(ep) (additional condition (2)).

$\Rightarrow$ P$_s$(Hit(v)) = s$_{ep}$(e).
$\Rightarrow$ s is uniform $\Rightarrow$ s$_{ep}$(e) = 1/ |Support$_s$(ep)|.

$\Rightarrow$ P$_s$(Hit(v)) = 1 / |Support$_s$(ep)|.

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Proof of Proposition 4 (cont.)

- Consider a vertex $v' \notin \text{Support}_s(vp)$.

  $\Rightarrow$ $\text{Support}_s(ep)$ is an Edge Cover of G (additional condition (1))
  $\Rightarrow$ there exists an edge $e \in \text{Support}_s(ep)$ such that $v' \in e$.

  $\Rightarrow$ $P_s(\text{Hit}(v')) \geq s_{ep}(e) = 1 / |\text{Support}_s(ep)|$ (s is uniform).

$\Rightarrow$ $P_s(\text{Hit}(v)) = \min_{v' \in V} P_s(\text{Hit}(v'))$, for each vertex $v \in \text{Support}_s(vp)$.

i.e. Condition (1).

- Condition (2): Similarly..

■

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Matching Nash Equilibria: Graph Theoretic Properties

- Proposition 5. In a Matching Nash equilibrium s,
  - $Support_s$(vp) is a Maximum Independent Set of G.
  - $Support_s$(ep) is a Minimum Edge Cover of G.

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

CONTRIBUTION 43

A characterization of Matching Nash Equilibria

- Theorem 25. A graph G admits a Matching Nash equilibrium if and only if G contains an Expanding Independent Set.

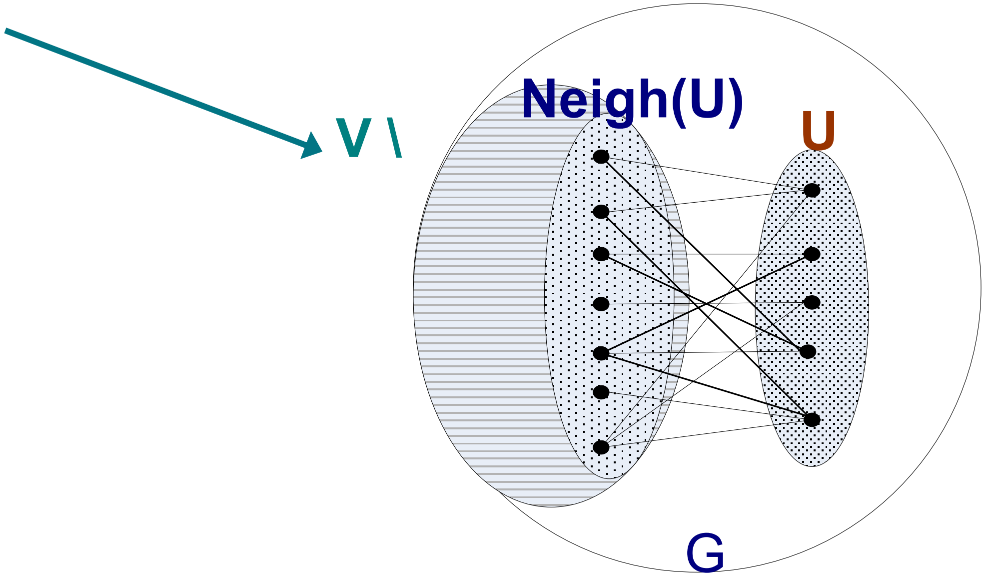

(Mavronicolas, Papadopoulou, Philippou & Spirakis (ISAAC 2005 & Algorithmica))

Game-Theoretic Characterizations of Graphs (1/3)

- Identified strong connections between Game Theory and Graph Theory. For example:
  - Discovered the first game-theoretic characterization of König-Egenváry graphs
  - Discovered game-theoretic analogs of graph-theoretic theorems.

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

Game-Theoretic Characterizations of Graphs (2/3)

- Theorem 25. The graph G admits a Matching Nash equilibrium if and it is König-Egenváry graph ($\alpha(G) = \beta'(G)$).

Proof.

- Assume that $\alpha(G) = \beta'(G)$
- IS = Max Independent Set
- EC = Min Edge Cover
- Construct a Uniform, Attackers Symmetric profile s with:
  - $Support_s(vp)$ = IS and $Support_s(ep)$ = EC.
- We prove that s is an Independent Covering profile

$\Rightarrow$ a Nash equilibrium.

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 46

Game-Theoretic Characterizations of Graphs (3/3)

- Assume now that G admits a Matching Nash equilibrium s.
  - By Proposition 5,

    $\Rightarrow$ $|\text{Support}_s(vp)| = |\text{Support}_s(ep)|$
  - by the definition of Matching Nash equilibria

    $\Rightarrow$ $\alpha(G) = \beta'(G)$. ■

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 47

Matching Nash Equilibria: Computation

- Since KÖNIG-EGENVÁRY MAX INDEPENDENT SET $\in \mathcal{P} \Rightarrow$

- Theorem 26. Finding a Matching Nash equilibrium can be solved in time $O\left(\sqrt{|V|}|E| \cdot \log_{|V|} \frac{|V|^2}{|E|}\right)$.

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

The Defense Ratio of Matching Nash Equilibria

- Proposition 5. In a Matching Nash equilibrium, the Defense Ratio is $\alpha(G)$.

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

Defender Uniform Nash Equilibria: A Characterization

- Theorem 27. A graph G admits a Defender Uniform Nash equilibrium if and only if there are non-empty sets $V' \subseteq V$ and $E' \subseteq E$ and an integer $r \geq 1$ such that:

(1/a) For each $v \in V'$, $d_{G(E')}(v) = r$.

(1/b) For each $v \in V \setminus V'$, $d_{G(E')}(v) \geq r$.

(2) V' can be partitioned into two disjoint sets V'_i and V'_r such that:

(2/a) For each $v \in V'_i$, for any $u \in Neigh_G(v)$, it holds that $u \notin V'$.

(2/b) The graph $\langle V'_r, Edges_G(V'_r) Å E' \rangle$ is an r-regular graph.

(2/c) The graph $\langle V'_I \cup (V \setminus V'), Edges_G V'_I \cup (V \setminus V')) Å E' \rangle$ is a $(V'_i, V \setminus V')$-bipartite graph.

(2/d) The graph $\langle V'_i \cup V \setminus V'), Edges_G(V'_i \cup V \setminus V') Å E' \rangle$ is a $(V \setminus V')$ - Expander graph.

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

Characterization of Defender Uniform Nash Equilibria: Illustration

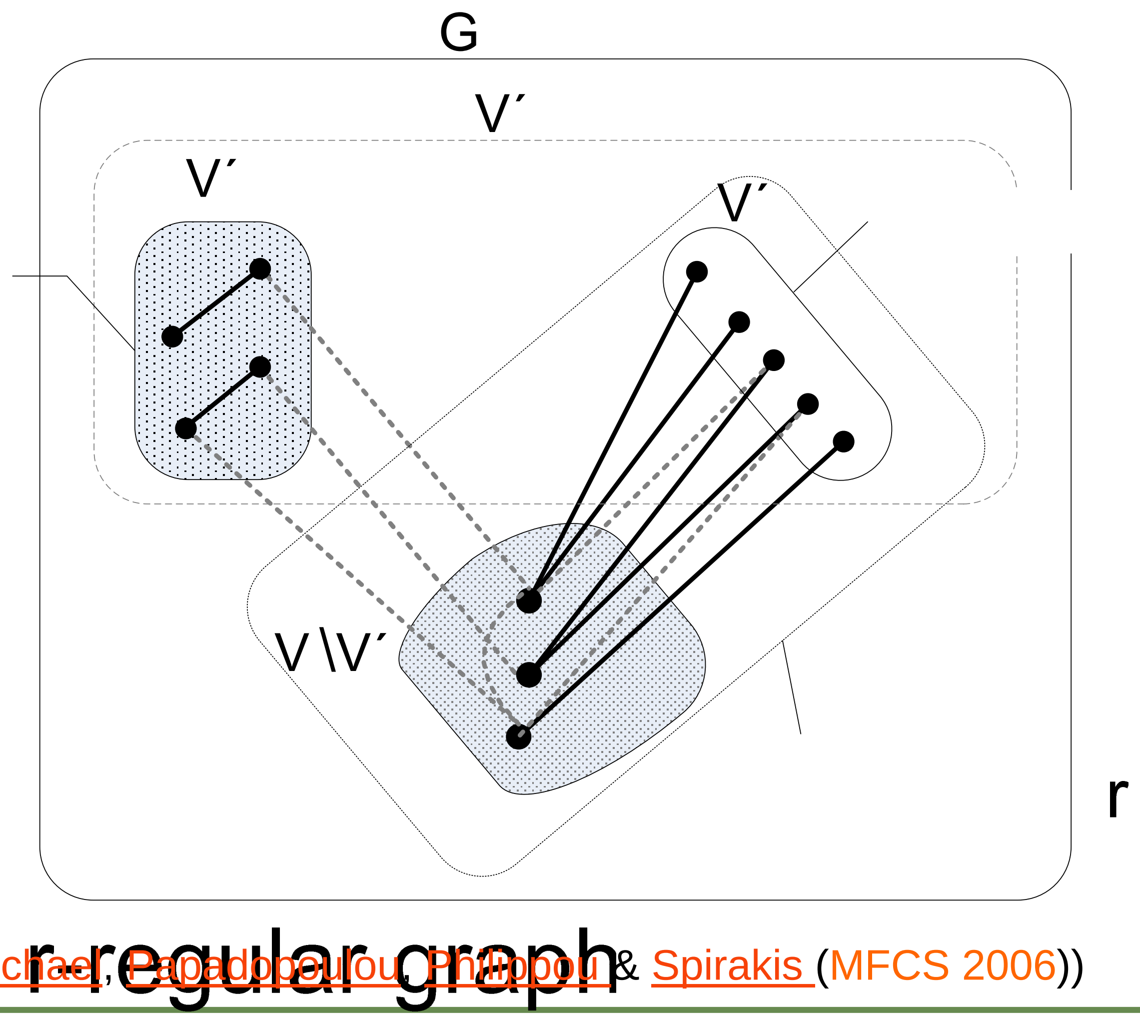

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 51

Defender Uniform Nash Equilibria: Complexity

- Theorem 28. The existence problem of Defender Uniform Nash equilibria is $\mathcal{NP}$-*complete.*

Proof.

- Reducing from
  - UNDIRECTED PARTITION INTO HAMILTONIAN CIRCUITS OF SIZE AT LEAST 6
    - proved to be $\mathcal{NP}$-complete.

■

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 52

The Defense Ratio of Defender Uniform Nash Equilibria

- Theorem 29. In a Defender Uniform Nash equilibrium, the Defense Ratio is $\left(\frac{\pi}{2}+1\right) \cdot |V|$, for some $0 \leq \pi \leq 1$.

(Mavronicolas, Michael, Papadopoulou, Philippou & Spirakis (MFCS 2006))

CONTRIBUTION 53

Bounds on the Price of Defense

- Theorem 30. The Price of Defense is $\geq \frac{|V|}{2}$.

- A graph is Defense-Optimal if it admits a Nash equilibrium with Price of Defense $= \frac{|V|}{2}$.

- Theorem 31. A graph is Defense-Optimal iff it admits a Fractional Perfect Matching.

(Mavronicolas, Papadopoulou, Persiano, Philippou & Spirakis (ICDCN 2006))

CONTRIBUTION 54

Fair Pricing Games

- Similar to the KP model, except that:
  - Weights are chosen according to some joint probability distribution D, which comes from a class Δ of probability distributions.
  - The Individual Cost to each agent choosing a resource is equal to

$$\text{Resource Cost} = \frac{\text{Total weight of agents choosing the resource}}{\text{Total number of agents choosing the resource}}$$

  - Social Cost = Maximum Resource Cost

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Strategies and Assignments

- A pure strategy for agent i is some specific resource.
- A mixed strategy for agent i is a probability distribution on the set of pure strategies.
- A pure assignment $\mathbf{L} \in M^n$ is a collection of pure strategies, one per agent.
- A mixed assignment $\mathbf{P} \in \mathrm{R}^{m \times n}$ is a collection of mixed strategies, one per agent.
  - i.e. p_i^j is the probability that agent i selects resource j.
  - The support of agent i is $S_i = \{ j \in M : p_i^j > 0 \}$.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

CONTRIBUTION 56

Resource Cost and Individual Cost

- Fix a pure assignment $\mathbf{L} = < l_1, l_2, \ldots, l_n >$.
- Resource demand on resource j : $W^j = \sum_{k \in N : l_k = j} w_k$.
- Resource congestion on resource j : $n^j = \sum_{k \in N : l_k = j} 1$.
- Resource Cost on resource j : $\mathrm{RC}^j = \frac{W^j}{n^j}$.
- Individual Cost for agent i : it is the Resource Cost of the resource she chooses, i.e. $\mathrm{IC}_i = \mathrm{RC}^{l_i} = \frac{W^{l_i}}{n^{l_i}}$.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

CONTRIBUTION 57

Expected Individual Cost

- Now fix a mixed assignment $\mathbf{P}$.
- The Conditional Expected Individual Cost IC_i^j of agent i on resource j is the conditional expectation of the Individual Cost of agent i had she been assigned to resource j.
- The Expected Individual Cost of agent i is

$$\mathrm{IC}_i = \sum_{j \in M} p_i^j \cdot \mathrm{IC}_i^j .$$

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Pure Nash Equilibria

- The pure assignment $\mathbf{L} = \langle l_1, l_2, \ldots, l_n \rangle$ is a pure Nash equilibrium if, for all agents i, the Individual Cost IC_i is minimized (given the pure strategies of the other agents).

- Thus, in a pure Nash equilibrium, no agent can unilaterally improve her own Individual Cost.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Mixed Nash Equilibria

- The mixed assignment **P** is a mixed Nash equilibrium if, for all agents i, the Expected Individual Cost IC_i is minimized (given the mixed strategies of the other agents), or equivalently, for all agents i,

$$\mathrm{IC}_i^j = \min_{k \in M} \mathrm{IC}_i^k \quad \forall j : p_i^j > 0$$

$$\mathrm{IC}_i^j \geq \min_{k \in M} \mathrm{IC}_i^k \quad \forall j : p_i^j = 0$$

- **P** is a fully mixed Nash equilibrium if additionally

$$p_i^j > 0 \quad \forall i \in N, \forall j \in M.$$

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

The Price of Anarchy

- Let $\mathbf{w}$ be a demand vector and $\mathbf{P}$ be a Nash equilibrium.
- The Social Cost is:

$$\mathrm{SC}(\mathbf{w},\mathbf{P}) = \mathrm{E}_{\mathbf{P}}\left(\max_{j\in M} \mathrm{RC}^{j}\right).$$

- Let $\mathbf{w}$ be a demand vector. The Optimum is:

$$\mathrm{OPT}(\mathbf{w}) = \min_{\mathbf{L}\in M^{n}} \max_{j\in M} \mathrm{RC}^{j}.$$

- The Price of Anarchy is:

$$\mathrm{PA} = \max_{\mathbf{w},\mathbf{P}} \frac{\mathrm{SC}(\mathbf{w},\mathbf{P})}{\mathrm{OPT}(\mathbf{w})}.$$

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

CONTRIBUTION 61

The Diffuse Price of Anarchy

- Assume demands are chosen according to some joint probability distribution D, which comes from some (known) class Δ of possible distributions.

- We define the Diffuse Price of Anarchy to be

$$\mathrm{DPA}_{\Delta} = \max_{D \in \Delta} \left(E_D \underbrace{\max_{NE} \frac{SC(NE)}{\min_{\alpha} SC(\alpha)}}_{\text{Price of Anarchy}} \right)$$

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Motivation

- The proposed cost mechanism is used in real life by:
  - Internet service providers
  - Operators in telecommunication networks
  - Restaurants offering an “all-you-can-eat” buffet
- The cost mechanism is **fair** since
  - No resource makes profit
  - Agents sharing the same resource are treated equally

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

The Optimum

- Proposition 6. For any demand vector **w**,

$$\mathrm{OPT}(\mathbf{w}) = \frac{W}{n}.$$

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Pure Nash Equilibria Inexistence

- Theorem 26. There is a pure Nash equilibrium if and only if all weights (demands) are identical.

Proof. (if)

- Let $w_i = w \;\; \forall i \in N.$
- Then, in any pure assignment **L**,

$$\mathrm{RC}^j = w \;\; \forall j \in M \quad \Rightarrow \quad \mathrm{IC}_i = w \;\; \forall i \in N.$$

- Hence any pure assignment is a pure Nash equilibrium.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Fully Mixed Nash Equilibria: Existence

- Theorem 27. There is always a fully mixed Nash equilibrium.

Proof.

- Let $\mathbf{F}$ be the fully mixed assignment with $f_i^j = \frac{1}{m} \; \forall i \in N, \forall j \in M.$
- In $\mathbf{F}$, and for all $i \in N$ and $j \in M$, it holds that

$$\mathrm{IC}_i^j = w_i \left(1 - \frac{1}{m}\right)^{n-1} + \sum_{k=2}^{n} \frac{1}{k} \left(\frac{1}{m}\right)^{k-1} \left(1 - \frac{1}{m}\right)^{n-k} \left(\binom{n-1}{k-1} w_i + \binom{n-2}{k-2} W_{-i} \right)$$

- i.e. the Conditional Expected Individual Cost of an agent i on resource j is independent of j, so $\mathbf{F}$ is a fully mixed NE. ■

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Fully Mixed Nash Equilibria: Uniqueness

- Theorem 28. The fully mixed Nash equilibrium **F** is the unique Nash equilibrium in the case of 2 agents with non-identical demands.

Proof.

- Consider an arbitrary Nash equilibrium **P**.
- Let S_1, S_2 be the support of agent 1, 2 respectively.
- W.l.o.g., assume that $w_1 > w_2$.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Fully Mixed Nash Equilibria: Uniqueness (cont.)

Proof. (continued)

- We can prove (by contradiction) that $S_1 = S_2 = M$.
- Now fix $j, k \in M$. Then

$$\mathrm{IC}_1^j = \mathrm{IC}_1^k \Leftrightarrow p_2^j = p_2^k \Leftrightarrow p_2^j = \frac{1}{m} \forall j \in M$$

$$\mathrm{IC}_2^j = \mathrm{IC}_2^k \Leftrightarrow p_1^j = p_1^k \Leftrightarrow p_1^j = \frac{1}{m} \forall j \in M.$$

- Hence P=F. ■

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

The Price of Anarchy: Lower Bound (1/2)

- Theorem 29. The Price of Anarchy is ˝έ $\frac{n}{2e}$.

Proof.

- First observe that $\mathrm{SC}(\mathbf{w},\mathbf{F}) \geq \left(\frac{1}{m}\right)^n \left(m(m-1)^{n-1} w_1\right)$
- Fix a demand vector $\mathbf{w}$ with $w_1 = \Theta\left(2^n\right)$ and $w_i = 1 \quad \forall i \neq 1$
- Then $\frac{w_1}{W} \geq \frac{1}{2}$.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

The Price of Anarchy: Lower Bound (2/2)

Proof. (continued)

- Now $\mathrm{PA} = \max_{\mathbf{w},\mathbf{P}} \left(\frac{n}{W} \cdot \mathrm{SC}(\mathbf{w},\mathbf{P}) \right)$

$$\geq \max_{\mathbf{w}} \left(\frac{n}{W} \cdot \mathrm{SC}(\mathbf{w},\mathbf{F}) \right)$$

$$\geq \max_{\mathbf{w}} \left(\frac{n w_1}{W} \cdot \left(\frac{m-1}{m} \right)^{n-1} \right)$$

$\geq \frac{n}{2e}$ for $m=n$, as needed. ■

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

CONTRIBUTION 70

The Price of Anarchy: Upper Bounds (1/2)

- Theorem 30. For two agents ($n = 2$), the Price of Anarchy, is $PA < 2 - \frac{1}{m}$.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

The Price of Anarchy: Upper Bounds (2/2)

- Theorem 31. The Price of Anarchy is $\text{PA} \leq \frac{n \cdot w_1}{W}$.

Proof.

- Fix any $\mathbf{w}$. For any pure assignment,

$$\frac{W^j}{n^j} \leq w_1 \quad \forall j \in M : n^j > 0$$

- Hence, for any Nash equilibrium $\mathbf{P}$,

$$\text{SC}(\mathbf{w}, \mathbf{P}) = \text{E}_{\mathbf{P}}\left(\max_j \frac{W^j}{n^j} \right) \leq w_1 \quad \Rightarrow \quad \text{PA} \leq \frac{n \cdot w_1}{W}.$$

■

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

CONTRIBUTION 72

The Diffuse Price of Anarchy (1/2)

- The class Δ of Bounded, Independent Probability Distributions:
  - Weights are independent, identically distributed random variables such that:
    - $\exists\, \delta_D < \infty \;:\; w_i \leq \delta_D$ for each $D \in \Delta$
    - $\exists\, l_\Delta > 0 \;:\; \dfrac{\delta_D}{E_D(w_i)} \leq l_D$ for each $D \in \Delta$

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

CONTRIBUTION 73

The Diffuse Price of Anarchy (2/2)

- Theorem 29. Consider the class Δ of bounded, independent probability distributions. Then,

1. $\mathrm{DPA}_{\Delta} \le \dfrac{\lambda_{\Delta}}{1-\lambda_{\Delta}\sqrt{1/2\ln n}} + n\exp\left(-\dfrac{n}{\ln n}\right)$

2. $\lim_{n\to\infty} \mathrm{DPA}_{\Delta} \le \lambda_{\Delta}$.

(Mavronicolas, Panagopoulou, & Spirakis (WINE 2005 & Algorithmica))

Research Plans

- Study further the complexity of Nash equilibria.
- Study further the evaluation of Nash equilibria via the Price of Anarchy.
- Apply game-theoretic modeling and analysis to problems from practical applications (e.g., caching, bandwidth allocation, security, network formation, etc.).
- Develop the field of Selfish Distributed Computing as a realistic reformulation of Fault-Tolerant Distributed Computing.
- Develop the field of Byzantine Game Theory.

More Concrete Research Plans

- Further develop the theory of complexity classes $\mathcal{PLS}$ and $\mathcal{PPAD}$ and classes in the Polynomial Time Hierarchy in relation to the problem of computing and counting equilibria.

  (Jointly with B. Monien and K. Wagner)

- Develop the algorithmic theory of games with collusion.

  (Jointly with F. Meyer auf der Heide)

- Study further security games with interdependencies.

  (Jointly with B. Monien, V. Papadopoulou, A. Philippou and P. Spirakis)

- Develop the algorithmic theory of tremble equilibria.

  (Jointly with P. Spirakis)

Research Funds

- European Union , IST Program:
  - ALCOM-FT
  - FLAGS (Global Computing 1) } completed
  - DELIS
  - AEOLUS (Global Computing 2) } In progress
- Cyprus Research Foundation, Funds for supporting research collaboration between two European countries:
  - EPDS (Cyprus-Greece) } completed  
  - ALGATHE (Cyprus-France) } in progress  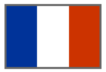

Thanks