A Study on the Stability and Efficiency of Graphical Games with Unbounded Treewidth

Anisse Ismaili
LIP6-UPMC, Paris, France
anisse.ismaili@lip6.fr

Evripidis Bampis
LIP6-UPMC, Paris, France
evripidis.bampis@lip6.fr

Nicolas Maudet
LIP6-UPMC, Paris, France
nicolas.maudet@lip6.fr

Patrice Perny
LIP6-UPMC, 4 place Jussieu,
75005 Paris, France
patrice.perny@lip6.fr

ABSTRACT

Graphical games (GG) provide compact representations of multiplayer games involving large populations of agents when influences among them have some locality property. The notion of pure Nash equilibrium (PNE), not requiring randomized strategies, is a fundamental stability concept. However, recent results show that for many natural topologies, a PNE is very unlikely to exist when the number of agents is large, which challenges the relevance of PNE in large GG.

In this paper, we investigate how far we can get from the notion of individual stability captured by the concept of PNE, by only requiring agents to be almost in best-response (ϵ-Nash), or by requiring almost all agents to be in best-response. We study these approximated notions of PNE for different topologies, including graphs with unbounded treewidth, like grids. This makes the problem computationnally very challenging and requires the comparison and use of several algorithmic solutions. Our results reveal surprisingly good asymptotic properties, tempering the claim that individual stability is not a relevant notion for large GG. Finally, as approximated PNE provide various tradeoffs between stability and social utility maximization, we propose an approach to construct a minimal-size e-covering of all feasible Pareto-dominant tradeoffs.

Categories and Subject Descriptors

I.2.11 [**Distributed Artificial Intelligence**]: Multiagent systems

General Terms

Algorithms, Theory, Experimentation

Keywords

Graphical Games, Approximation

Appears in: *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2013), Ito, Jonker, Gini, and Shehory (eds.), May, 6–10, 2013, Saint Paul, Minnesota, USA.*

1. INTRODUCTION

Games involving a large number of players require a prohibitive amount of space in order to be represented in normal form. However, in many cases, the payoff functions of agents do not depend on all other agents. Graphical games (GG) provide compact representations of multiplayer games involving large populations of agents when influences among them exhibit such locality property [14, 13]. In a GG, each vertex is an agent, and an agent's payoff is a function of his action and the actions of the agents in the inner-neighborhood.

This paper investigates the stability and efficiency of graphical games, with agents playing pure strategies (*i.e.* not randomized). In this context, the notion of Pure Nash Equilibria (PNE) is arguably the most fundamental stability notion: no agent has an individual incentive to deviate. Stability usually comes at the price of efficiency: the ratio of the best (resp. worst) NE over the social welfare optimum is known as the price of stability [1] (resp. price of anarchy).

Unfortunately, there is no guarantee for such a PNE to always exist (in contrast to the case of mixed Nash equilibria). The existence problem of PNE in general GG has been known to be hard for a long time [10].

By exploiting non-serial dynamic programming [2] as used in graphical models for probabilistic inference, [5] provides an algorithm which handles efficiently graph structures with bounded (or logarithmic, when the number of actions is bounded) treewidth. One can also cite [19] which assumes a bounded treewidth too. Recently, [12] even showed that this criteria suffices to fully characterize hard *vs.* easy instances in GG. Essentially, GG are easy if and only if they have a bounded treewidth (after their sinks have been iteratively removed). The case of unbounded treewidth is nevertheless of great interest, since many natural graphs fall in this category (e.g. grids).

Another important issue when there are several possible PNE is to search for the best one in terms of some welfare function [3]. However, recent results show that when the number of agents becomes large, the likelihood to find at least one PNE converges to 0 for many graph topologies. Indeed, in [7] different natural topologies are investigated, both theoretically and experimentally. They show that the structure heavily affects the probability of existence of a PNE: for instance, while the probability converges to 0 in the case of trees with an unbounded diameter, it remains as high as $1 - \frac{1}{e}$ for the standard (random) normal form, or bipartite graphs. They suggest that the probability of existence increases as the length of paths diminishes. From the results of [7], there is no evidence that the probability of PNE existence in general graph topologies tends to a strictly positive value when the diameter diverges more quickly than $\Omega(\log(n))$. The same problem is addressed in [4] for graphs drawn from the Erdös-Renyi $\mathcal{G}(n, p)$ model (where n is the number of vertices and $p \in [0, 1]$ is the probability for an unoriented edge to exist). They essentially show that when $\frac{c}{n^2} < p < \frac{\ln(n)}{2n}$ (medium connectivity) the probability for a PNE to exist converges to 0, and when $p > 2\frac{\ln(n)}{n}$ (high connectivity) the number of PNE**s** converges to a poisson(1) distribution (the probability of existence is again $1 - \frac{1}{e}$).

These results may challenge the relevance of the notion in GG involving a large number of agents (at least for the aforementioned topologies). However, they provide a partial picture of the problem. In particular, they mostly pay attention to the 2-action case and remain constrained by the exact definition of PNE. In particular, it does not address the question of whether it is possible to find a profile of strategies which is "almost" a PNE. This is the first problem that we study in this paper. This requires to define what an approximate PNE is. As a PNE requires *all* agents to *strictly* play a best-response, we shall consider two possible approximations in this paper. The first one is to modify the condition under which agents may deviate. Specifically, we consider the so-called multiplicative ϵ-Nash equilibrium. The rationale of this notion is that the relative incentive to deviate for agents may be small in practice. The second one that we call k-Nash consists in relaxing the constraint that all agents are in best-response. The rationale of this approximation is that it may only require, for the system designer, to "convince" or "impose" the strategy of a small number of agents. This notion bears some similarity with the notion of Stackelberg threshold [17], with the difference that this measures the number of agents whose strategy must be fixed in order to guarantee a social optimum at equilibrium, while in our case, we seek to guarantee that an equilibrium exists. We also note that the study of k-Nash approximations is very fundamental, since it does not rely on the payoff random distributions, but just on best response tables.

Unfortunately, as previously mentioned, instances involving unbounded treewidth are computationally very challenging to handle. We deploy a range of algorithmic solutions (junction tree algorithms [5], a SAT formulation due to [7], and Mixed Integer Linear Programming), and compare their respective efficiency on the different problems we face. This allows us to perform different experiments involving a significant number of agents (typically at least 100 agents) and action-sets of size greater than 2. Our results shed a new light on the PNE problem in large GG: in a nutshell, we show that in various GG whose PNE probability of existence converges to 0, we actually get very close to PNE (in particular, a very small *constant* fraction of agents may be unsatisfied), and that the action-set size may greatly affect the results (in particular, it contradicts the assumption that the divergence of the graph diameter may be a sufficient condition to conclude on the asymptotic inexistence of PNE).

We next turn our attention to the tradeoff between stability and utilitarian optimality. Equipped with a suitable approximation of PNE, it is actually possible to generalize the notion of price of stability by seeing it as a bi-criteria problem. The set of non-dominated tradeoffs would provide precious indications for the system supervisor ("what gain in social welfare can be expected by relaxing this much of the stability requirement?"). Specifically, we question how the price of stability evolves as a function of the best response ϵ-relaxation, by computing the values of an $(1 + e)$-approximated bi-criteria Pareto-set (ϵ, v) of minimal size, with a minimal number of calls to an utilitarian optimization problem and an ϵ-stability optimization problem, thanks to a greedy algorithm defined in [6].

The remainder of this paper is as follows. In Section 2 we introduce the necessary background on graphical games. In Section 3, we provide the different algorithmic approaches we use, and discuss how they perform for the different problems we have to deal with. Approximated notions of PNE are introduced and investigated in Section 4. Finally, Section 5 details the methodology to obtain an approximated Pareto-set for the stability vs. efficiency tradeoff.

2. PRELIMINARIES

We remind common notations of multiagent games. The set of agents is $N = \{1, \ldots, i, \ldots, n\}$. Each agent i chooses his action a_i in his particular action-set A_i. An action-profile is any element: $(a_1, \ldots, a_i, \ldots, a_n) \in \prod_{i=1}^n A_i$ denoted by $a \in A$. An agent i's adversary-action-profile is denoted $a_{-i} \in \prod_{j \neq i} A_j$ and his payoff function is $v_i : A_i \times A_{-i} \to \mathbb{R}$.

Definition 1 *A multiagent game (or n-agent game when $|N| = n$) is a t-uple $\Gamma = \left(N, \{A_i\}_{i \in N}, \{v_i\}_{i \in N}\right)$.*

Nash equilibria formulate the "individual stability" of an action-profile $(a_i, a_{-i}) \in A$: each agent i, given the fixed adversary-action-profile a_{-i}, has no incentive to change unilaterally his action a_i, to get an individually prefered action-profile.

Definition 2 *Given a multiagent game Γ and a player i, a best response function $BR_i : A_{-i} \to \mathcal{P}(A_i)$ maps adversary-action-profiles to the set of actions which "satisfy" i.*

Definition 3 *An action-profile $a \in A$ is a Nash equilibrium if and only if for every agent i, $a_i \in BR_i(a_{-i})$.*

If we consider $BR_i(a_{-i}) = \text{argmax}_{a_i \in A_i}\{v_i(a_i, a_{-i})\}$, then we are working on Pure Nash Equilibria (PNE). If we consider $BR_i^\epsilon(a_{-i}) = \{a_i \in A_i | \forall a_i' \in A_i : \frac{v_i(a_i', a_{-i})}{v_i(a_i, a_{-i})} \leq (1 + \epsilon)\}$ (assuming positive utilities) , then we are working on multiplicative ϵ-approximated Nash Equilibria (ϵ-PNE).

In the most general normal forms, each payoff function requires $O\left(\prod_{i \in N} |A_i|\right)$ space, which is not compact if the number of agents is unbounded. A graphical game (GG) catches the locality of payoff functions. Its support graph is an oriented graph $G = (N, E)$, where N is the set of agents, E are the edges which define dependencies between payoffs, and agent i's inner-neighborhood (including himself) is denoted $\nu(i)$. Player i's payoff function depends only on the actions of players in $\nu(i)$. A local-action-profile is any element: $(a_i, \ldots) \in \prod_{j \in \nu(i)} A_j$ denoted by: $a^{(i)} \in A^{(i)}$. In a graphical game, agent i's payoff is a function $v_i : A^{(i)} \to \mathbb{R}$, which by the way requires $O\left(\prod_{j \in \nu(i)} |A_j|\right)$ space. Hence GG are a compact multiagent games representation when $\Delta = \max_{i \in N} \{|\nu(i)|\}$ is bounded and small.

Definition 4 *A graphical game (GG) with support graph* $G = (N, E)$ *is a t-uple* $\Gamma = \left(G, \{A_i\}_{i\in N}, \{v_i\}_{i\in N}\right)$.

What we are mostly going to consider are unoriented graphical games, where $(i, j) \in E \Leftrightarrow (j, i) \in E$ (payoff dependencies between agents are mutual).

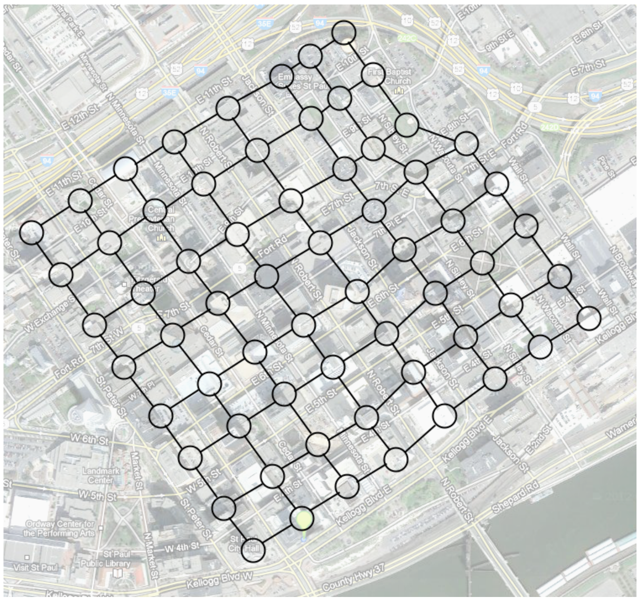

Figure 1: St Paul Town-Game

Example 1 *Consider the example of Figure 1 inspired by the Road-Game [20]. In the St Paul Town-Game, each agent possesses one block of land in a square-grid-town of size* $n = L \times L$*. Each agent must choose what to build on his block of land: a garden, a residential complex, a factory, or a shopping mall. Clearly, the payoff of a building depends on what is build on the four adjacent blocks of land. Unfortunately, square-grid-graphs of size* $n = L \times L$ *have unbounded treewidth* L*, and an even higher hypertreewidth.*

We might also want to find an action-profile which maximizes the average payoff, like in the VCG mechanism and Clarke's prices.

Definition 5 *Given a GG* $\Gamma = \left(G, \{A_i\}_{i\in N}, \{v_i\}_{i\in N}\right)$ *the utilitarian optimization problem (denoted as* $\max \sum$*-GG) is to find* $a^* \in argmax_{a\in A}\{\sum_{i\in N} v_i(a^{(i)})\}$.

Given a fixed support graph $G = (N, E)$, we study *random-payoff graphical games* with respect to the following distribution: each payoff values $v_i(a^{(i)})$ of the functions $v_i : A^{(i)} \rightarrow \mathbb{R}$ are drawn uniformly and independently from the interval $[0, 1]$. For the PNE and the k-Nash problems, it is equivalent to suppose that each adversary-local-action-profile has a singleton best-response. But for the ϵ-Nash problem, it is an important (but common) hypothesis that can influence the results.

Throughout this paper, $B^{(i)} = \{a^{(i)} \in A^{(i)} | a_i \in BR_i\left(a_{-i}\right)\}$ will denote the best-responses local-action-profiles, and $\overline{B}^{(i)} = A^{(i)} \setminus B^{(i)}$ the non-best-responses local-action-profiles. For ϵ-best-responses, $\epsilon' < \epsilon \Rightarrow \overline{B}^{(i)}_{\epsilon'} \supset \overline{B}^{(i)}_{\epsilon}$.

3. COMPUTATION IN GRAPHICAL GAMES

In this section, we present various ways to compute a PNE in a *general* GG. Recall that this is an NP-Hard problem. It is known that this problem's structure is close to a logical problem [7]. In what follows, we first remind how junction trees can be used when (hyper)treewidth is bounded. Then we detail the SAT formulation that [7] briefly proposes and massively uses. We also propose a MIP formulation of the same problem, and we show how this MIP can be adapted to the case of an unbounded maximal degree Δ, with a cutting plane algorithm polynomial in Δ, for non-best-responses.

Junction Tree Algorithm. Computing a PNE in a GG $\Gamma = \left(G, \{A_i\}_{i\in N}, \{v_i\}_{i\in N}\right)$ can be reduced in various ways to the Junction Tree algorithm. For instance, [5] builds a maximum likelihood problem in a Markov random field (ML-MRF) over $A_1 \times \ldots \times A_n$ which is clique-factorized and whose cliques roughly correspond to the agents neighborhood: $C_i = \nu(i)$ and $\mathbb{P}(a) = \prod_{i\in N} \phi_{C_i}(a^{(i)})$. The potential functions are

$$\phi_{C_i}(a^{(i)}) = \begin{cases} 1 & \text{if} \quad a^{(i)} \in B^{(i)} \\ \epsilon & \text{otherwise} \end{cases}$$

in such a way that the maximum likelihood are 1 if and only if there is a PNE. With the same reduction, one can prove the NP-Hardness of $\max \sum$-GG. It is sufficient to change the multiplicative algebra into an additive algebra.

One can remark that if we choose $\epsilon = e^{-1} = \exp(-1)$, this ML-MRF not only decides on the existence of a PNE, but also maximizes the number of agents in a best-response local-action-profile. Indeed, if π^* is the maximal log-likelihood of this ML-MRF, then $-\pi^*$ is precisely the minimal number of unsatisfied agents. Unfortunately, these algorithms assume a bounded degree (or at most logarithmic with a fixed number of actions) treewidth, for tractability. They behave very badly in the general case, including the St Paul Town-Game.

SAT. Let us now introduce a decision variable $x_{i,a_i} \in \{0, 1\}$ for each agent i and each one of his action a_i. These variables can be both interpreted as a boolean or a binary variable, whose meaning is "agent i chooses action a_i". From this, given a GG $\Gamma = \left(G, \{A_i\}_{i\in N}, \{v_i\}_{i\in N}\right)$, the SAT formulation proposed in [7] is obtained by two sets of clauses. The first one ensures that each agent chooses one and only one action:

$$\forall i, \bigvee_{a_i\in A_i} x_{i,a_i}$$

$$\forall i, \forall\{a_i, a'_i\} \subset A_i, \neg x_{i,a_i} \vee \neg x_{i,a'_i}$$

The second set of clauses forbids the non-best-responses local-action-profiles :

$$\forall i, \forall a^{(i)} \in \overline{B}^{(i)}, \bigvee_{j\in\nu(i)} \neg x_{j,a^{(i)}_j}$$

One can remark that, this formulation can also be used to answer for a fixed ϵ the decision problem "Is there an ϵ-Nash Equilibrium?", by simply prohibiting $\overline{B}^{(i)}_{\epsilon}$.

MIP. For the first MIP formulation that we propose, we use the same variables and the same constraints, and we introduce the relaxation variables $\rho_i \in [0, 1]$ the sum of which must be minimized so that a PNE exists if and only if the optimal value is 0. First of all we need to add constraints enforcing that i chooses exactly one action:

$$\forall i \in N, \forall a_i \in A_i : \quad x_{i,a_i} \in \{0,1\} \quad (1)$$
$$\forall i \in N : \quad \sum_{a_i \in A_i} x_{i,a_i} = 1 \quad (2)$$

Then we impose that for each non-best-response local-action-profile, at least one decision of the inner-neighborhood $\nu(i)$ is 0. These contraints are relaxed by using continuous variables ρ_i.

$$\forall i \in N : \quad \rho_i \in [0,1] \quad (3)$$
$$\forall i \in N, \forall a^{(i)} \in \overline{B}^{(i)} : \quad \sum_{j \in \nu(i)} x_{j,a_j^{(i)}} \leq |\nu(i)| - 1 + \rho_i \quad (4)$$

We can then write the MIP addressing the PNE problem in GG:

$$\begin{array}{ll} \text{minimize} & \sum_{i=1}^{n} \rho_i \\ \text{subject to} & (1) \wedge (2) \wedge (3) \wedge (4) \end{array}$$

By introducing a cutoff if $\max_i \rho_i > 0$, we avoid many exploration nodes that are useless for the Pure NE problem. The number of non-best-response local-action-profile constraints is an exponential of Δ. Hence this is a compact formulation if Δ is bounded.

Techniques. We tested three computation techniques: maximum likelihood in a Markov random field with the junction tree algorithm implemented in libDAI [15], the SAT formulation with the miniSAT solver [8], and the MIP formulation in GuRoBi [11]. We shall summarize how these approaches fit the different problems considered in the following section.

Figure 2 compares the computation times in unbounded (hyper)treewidth GGs, using our MIP formulation (ILP) and junction-trees (ML_MRF). Both techniques are surpassed by miniSAT (which required an almost 0 time, here).

Figure 2: Computation average times for the PNE problem in square grids made of 2 actions agents.

4. APPROXIMATING PURE NE

In this section we investigate two different kinds of approximation of pure equilibria in GG. We conduct experiments with these approximations on support graphs whose random payoff GG are known to be unfortunately not likely to have a PNE when n is high: for paths, binary trees [7], and medium connected Erdös-Renyi [4], where we know that $\mathbb{P}(\exists a : \text{ PNE }) \to 0$, when $n \to \infty$, and for square grids, for which we experimentally measured (see Figure 4) that the PNE probability tends to 0 for small action-sets.

ϵ-Nash. The first notion we study is the multiplicative ϵ approximated NE, which minimizes $\epsilon(a) = \max_{i \in N} \{\epsilon_i(a^{(i)})\}$ (where $\epsilon_i(a^{(i)})$ is defined by $(1 + \epsilon_i(a^{(i)}))v_i(a^{(i)}) = \beta_i(a^{(i)})$ and $\beta_i(a^{(i)})$ is the best individual value that i can achieve with an individual deviation). This general idea of this approximation is standard in the literature (although we note that it is usually used in the additive sense), and reflects the intuition that "agents are indifferent to sufficiently small gains" [18]. We emphasize that this notion is often used as a discretization to facilitate the computation of mixed equilibria. In this case, a common criticism is to say that it is not guaranteed to be close to an exact equilibrium (which does exist for sure). This argument does not hold here, as our objective is to minimize ϵ: when a PNE does exist, we indeed get $\epsilon = 0$, otherwise we find an action profile which is the less sensible to individual deviation: the smallest ϵ such that there is an action profile where agents are not likely to individually deviate when the gain is less than $(1 + \epsilon)$ times what they already have.

k-Nash. The second kind of approximation is to find an action-profile, which minimizes the number k of agents not playing a pure best-response (k-Nash), and hence to find the minimal fraction of unsatisfied agents that it is necessary to "convince" or "constrain" to obtain stability. This notion has a great advantage over the ϵ-Nash approximation: it does not rely on the payoff distributions, but just on singleton best response tables. We note that, in principle, the fraction of unsatisfied agents may converge to 0 even though the probability of PNE does so.

Methodology. Remark that both the maximum log likelihood and the MIP presented in Section 3 can be directly adapted to k-Nash: maximizing the number of agents playing a best response. If the support graph is of polynomial size and has a bounded treewidth, then one should use a junction tree algorithm. Otherwise, when the treewidth is unbounded, or when Δ is unbounded, one should use the MIP formulation when addressing k-Nash.

Now for the computation of ϵ-Nash, since for a fixed ϵ, the SAT formulation is the fastest one by a huge margin (regardless of the treewidth or the support graph's maximal degree Δ), it turns out to be faster than MIP to make a dichotomy over ϵ with the SAT formulation at each iteration.

We deduced the proportions of instances having a PNE from the fastest previous approximations. For the convenience of the reader, we summarize the characteristics of the graphs and the algorithmic approaches employed (for the different problems) in Table 1.

Results on paths and binary trees. We know from [7] that the probabilities of PNE to exist in paths and binary trees tend to 0 (they actually reach 0 with about 100 agents). Our experimental results are presented in Figure 3. We observe that these results are insensitive to the size of the action-set. On these structures, more actions for agents do not help to reach a PNE. Now regarding the approximations of PNE, one may intuitively expect the situation to worsen as the probability of PNE decreases. This is not the case: in both topologies, the k-Nash equilibria we found reveal that the average minimal fractions of unsatisfied agents tend to small constants, between 3% and 6%. It means that the average number of unsatisfied agents is strictly proportional to the total number of agents. Moreover, this convergence is very fast: even when the probability of PNE is as high as 50% (path, 2 actions), the average fraction of unsatisfied agents is already at 6%. In this sense, we are not far from individual rationality equilibrium (even though the number of

Graph parameters			Algorithmic approach	
Name	*Treewidth*	*Diameter*	*k-Nash*	*ϵ-Nash*
Path	1	n	Junction trees [1500]	Dichotomy of SAT [500]
Binary tree	1	$2 \log n$	Junction trees [1500]	Dichotomy of SAT [500]
Grid	$\sqrt{n}$	$2\sqrt{n}$	MIP [500]	Dichotomy of SAT [1500]
medium connect. Erdös-Renyi	$\Omega(\beta n)$, with $\beta > 0$ [9]	$\log n$	MIP [500]	Dichotomy of SAT [500]

Table 1: Summary of graph parameters and algorithmic solution for the different problems. The number in brackets indicates the number of payoffs GG per graph and action-size

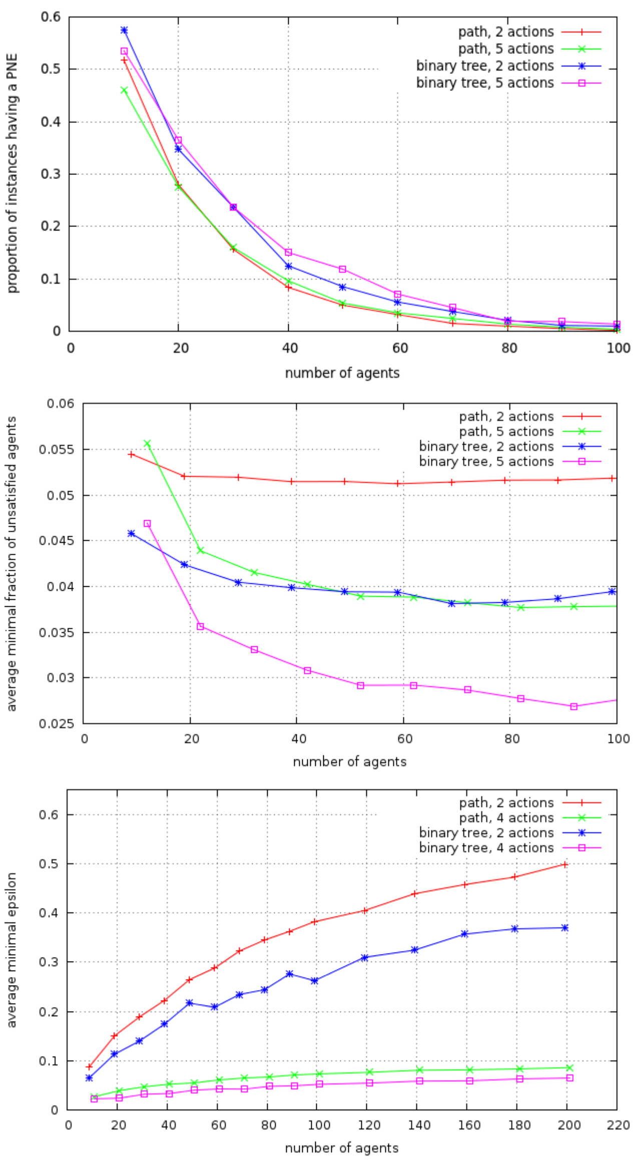

Figure 3: Indicators of PNE proximity in random-payoff GG : path and binary trees

unsatisfied agents grows, of course): in a path of 200 agents, the system designer may only have to fix the strategy of about 10 agents to obtain a stable situation.

Unlike for k-Nash, the experiments for ϵ-Nash show that in paths and trees, the results are affected by the size of the action-set. More precisely, the smaller the action-sets are, the harder it is to guarantee a small ϵ incentive for all agents in the same action-profile. To conclude, although more actions available for agents do not increase the probability of PNE, they do increase significantly the quality of the ϵ-Nash approximation (compare 200 agents on a path: ϵ is less than 10% for 4 actions *vs.* 50% for 2 actions).

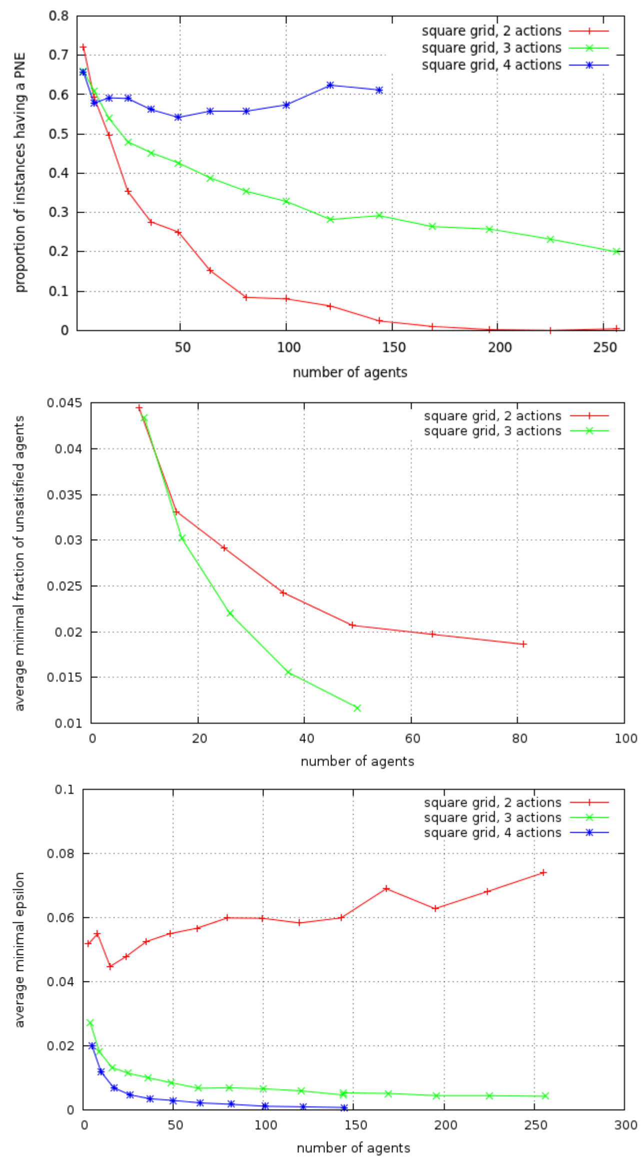

Figure 4: Indicators of PNE proximity in random-payoff GG : square grid graph

Results on grids. These results are in Figure 4. The probability of PNE existence is greatly affected by the action-set size in this topology (recall this was *not* the case for paths and binary trees). For 2 or 3 actions for each agent, these probabilities tend to 0 (at different rate) as the grid sizes increase. With action-sets of size 4, the influence of the action-set size can be observed even more dramatically: the probabilities of pure NE to exist do not tend to 0 any longer (they seem to tend to the $1 - \frac{1}{e}$ value). *This shows that the diameter of the graph cannot be a sufficient parameter to conclude on the fact that the probability of PNE existence converges asymptotically to 0 or not.* To the best of our knowledge, this is the first time that this is observed.

Turning our attention to k-Nash, we see that we are not so far from individual stability since the minimal fraction of unsatisfied agents are below 2%. About ϵ-Nash, ϵ tends to 0 with 3 or 4 actions, but does not with 2. We also note that this topology allows to reach very close to stable ϵ-Nash approximations: for instance, ϵ remains negligible for 250 agents with a 3-action set, when the probability of PNE has already decreased to 20%.

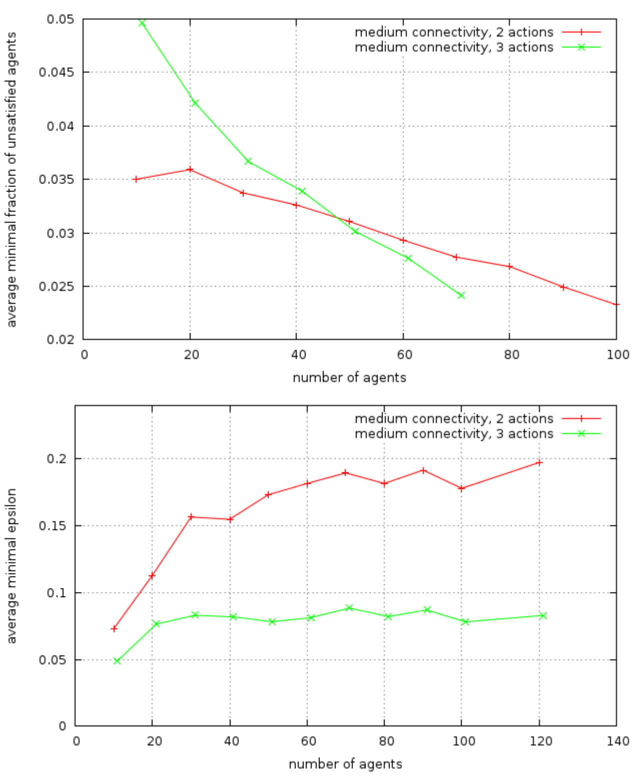

Figure 5: Indicators of PNE proximity in random-payoff GG: medium connectivity Erdös Renyi graph

Results on medium connectivity Erdös-Renyi. In medium connectivity ($\frac{c}{n^2} < p < \frac{\ln(n)}{2n}$) Erdös-Renyi graphs, we see that the convergence to 0 of PNE existence is not affected by the action-set size (not shown here). We can also observe that the average minimal fraction of unsatisfied agents decreases below 2%. As for the ϵ-Nash, the values obtained show that it is more difficult than with grids to get a close ϵ-approximation of PNE. For completeness, we mention that we also performed some experiments on small worlds (Watts-Stogratz) graphs, which are free-scale networks capturing some properties of biological, social, and computer networks. To sum up, we observed that the action-set size has a strong influence on both the PNE probability and the minimal ϵ values. By focusing our attention to $p \in [0.03, 0.375]$, we observe that we obtain very small ϵ values (from 3% to less than 1%). Note that when $p = 0.375$ the PNE probability has reached the $1 - \frac{1}{e}$ value [7].

Conclusions. Some general conclusions can be drawn from these experiments. First of all, even though the PNE probability tends to 0, in general, the approximated equilibria are very good. This is in line with some observations of [20] who mentioned that "it appears that, with random payoffs, fairly good equilibria often exist in pure strategies" after some experiments on grids (with 3-actions games, and under the additive ϵ-Nash interpretation). However, we also see that the picture is more complex. While k-Nash approximations always require only a small (decreasing or constant) fraction of agents to be convinced, it may sometimes be difficult to get small ϵ-approximations when the action-set is 2. But augmenting the size of the action-set always improves this approximation. As approximated PNE may provide various solutions, this gives us the flexibility to optimize the efficiency (in the utilitarian sense).

5. STABILITY VS SOCIAL UTILITY

Given a GG, we investigate the bi-criteria problem that consists in both maximizing "individual stability" and "overall utility". Formally, concerning stability, we will focus on the maximization of a strictly decreasing function of ϵ, $\varphi(a) = 1/(1 + \epsilon(a))$, which corresponds to the guaranteed fraction of individual rationality $\varphi_i(a^{(i)}) = v_i(a^{(i)})/\beta_i(a^{(i)})$ where $\beta_i(a^{(i)})$ is the best individual value that i can achieve with an individual deviation. Note that maximizing this function means that we minimize $\epsilon(a)$, which represents the worst relative opportunity loss among agents for action-profile a. Concerning utility, we maximize the standard utilitarian criterion defined by $v(a) = \sum_{i \in N} v_i(a^{(i)})$. Hence, to any a in A is assigned a performance vector $(\varphi(a), v(a))$ in the bi-criteria valuation space. The set of feasible performance vectors is then defined by $F(A) = \{(\varphi(a), v(a)), a \in A\}$ and the components are assumed to be bounded.

Since the two objectives considered are possibly conflicting, we are interested in finding action-profiles achieving different tradeoffs. To this end, determining a representative subset of feasible Pareto-optimal tradeoffs within $F(A)$ would provide useful information for the system supervisor. We recall that a vector $(\varphi, v) \in F(A)$ is Pareto-optimal if we cannot improve one dimension without downgrading the other. Note that the set $\mathcal{P}$ of Pareto-optimal vectors in $F(A)$ possibly includes a huge number of elements. Not only would its exact determination induce a prohibitive computational time, it would also provide the supervisor with an unnecessary large sample of possibilities. We propose instead to determine an approximation of $\mathcal{P}$ using the notions of e-dominance and e-covering.

e-dominance. For any $e = (e_\varphi, e_v) \in \mathbb{R}^2_+$ the e-dominance relation is defined on performance vectors $x = (x_\varphi, x_v), y = (y_\varphi, y_v) \in F(A)$ by $x \succcurlyeq_e y \Leftrightarrow [(x_\varphi(1 + e_\varphi) \geq y_\varphi)$ and $(x_v(1 + e_v) \geq y_v)]$.

e-covering. For any $e = (e_\varphi, e_v) \in \mathbb{R}^2_+$ and any set $S \subseteq$

$F(A)$, the set $T \subseteq S$ is said to be an e-covering of S if $\forall s \in S, \exists t \in T, t \succcurlyeq_e s$.

In general, multiple e-covering of the Pareto set $\mathcal{P}$ exist. The set $\mathcal{P}$ itself in an obvious example of e-covering of $\mathcal{P}$. More interestingly, the existence of e-covering of polynomial size is established for general multiobjective combinatorial problems with a bounded number of criteria [16]. Here, for any given $e = (e_\varphi, e_v) \in \mathbb{R}^2_+$ we would like to be able to determine an e-covering of $\mathcal{P}$ as small as possible, ideally an e-covering of $\mathcal{P}$ of minimal cardinality. Since we have only two objectives we can resort to a greedy approach proposed in [6] for computing an e-covering of minimal cardinality in general bi-criteria optimization problems. Let us explain how this approach can be further specified to explore the feasible tradeoffs attached to the action-profiles in a GG.

The construction proposed in [6] relies on solving a sequence of optimization problems alternating two complementary subproblems:

Restrict-v(α). For any given value α, we want to maximize v subject to the constraint $\varphi \geq \alpha$, or answer *no* when no such solution exists.

Restrict-φ(γ). For any given value γ, we want to maximize φ subject to the constraint $v \geq \gamma$, or answer *no* when no such solution exists.

We present now how we can solve these two subproblems in our context. Since these two problems are NP-hard, and since they include too much numerical information to resort to SAT reformulation, we focus on Mixed Integer Programming formulations (MIP).

Solving Restrict-v(α). With no loss of generality, we assume that payoffs are positive ($v_i(a) \geq 0$). If it is not the case, a simple translation of payoffs is sufficient. We remind players action variables $x_{i,a_i} \in \{0,1\}$ whose value is 1 iff player 1 chooses action a_i. First of all we need to add constraints enforcing that i chooses exactly one action :

$$\forall i \in N, \forall a_i \in A_i : \quad x_{i,a_i} \in \{0,1\} \tag{5}$$

$$\forall i \in N : \quad \textstyle\sum_{a_i \in A_i} x_{i,a_i} = 1 \tag{6}$$

Then we define the variables of local action profiles $z_{i,a^{(i)}}$, in such a way that we have $z_{i,a^{(i)}} = \bigwedge_{j \in \nu(i)} x_{j,a_j^{(i)}}$:

$$\forall i \in N, \forall a^{(i)} \in A^{(i)} : \quad 0 \leq z_{i,a^{(i)}} \leq 1 \tag{7}$$

$$\forall i \in N, \forall a^{(i)} \in A^{(i)}, \forall j \in \nu(i) : \quad z_{i,a^{(i)}} \leq x_{j,a_j^{(i)}} \tag{8}$$

And make players payoffs v_i consistent with the activated local action profile:

$$\forall i \in N : \quad v_i = \textstyle\sum_{a^{(i)} \in A^{(i)}} v_i(a^{(i)}) z_{i,a^{(i)}} \tag{9}$$

To complete the formulation of Restrict-v(α), we must also require that $\varphi \geq \alpha$:

$$\forall i \in N, \forall a^{(i)} \in \overline{B}^{(i)}_\alpha, z_{i,a^{(i)}} = 0 \tag{10}$$

To avoid very large integrality gaps, the following reinforcement is useful:

$$\forall i \in N, \sum_{a^{(i)} \in A^{(i)}} z_{i,a^{(i)}} = 1 \tag{11}$$

Moreover constraints (11) become necessary when the problem is empty due to constraints (10). We can then write our MIP formulation solving Restrict-v(α):

$$\begin{array}{ll} \text{maximize} & \sum_{i=1}^n v_i \\ \text{subject to} & (5) \wedge (6) \wedge (7) \wedge (8) \wedge (9) \wedge (10) \wedge (11) \end{array}$$

Since all $v_i(a^{(i)})$ are positive, and we maximize, the local action profiles variables are consistent.

Solving Restrict-φ(γ). In this problem, we want to maximize $\varphi(a) = 1/(1+\epsilon(a))$ with $\epsilon(a) = \max_{i \in N}\{\epsilon_i(a)\}$ and $\epsilon_i(a) = \sum_{a^{(i)} \in A^{(i)}} \epsilon_i(a^{(i)}) z_{i,a^{(i)}}$. We recall that our actual objective is to minimize $\epsilon_i, i \in N$ under the constraint that the total payoff is at least γ. However, it is not possible to proceed exactly as above to formulate the restrict-φ problem using variables $z_{i,a^{(i)}}$ because now the linearizations as well as the consistencies of $z_{i,a^{(i)}} = \bigwedge_{j \in \nu(i)} x_{j,a_j^{(i)}}$ are lost. To recover the necessary monotonicity of the objective function with respect to $z_{i,a^{(i)}}$ variables we maximize a strictly decreasing function of ϵ. We therefore define $\varphi_i(a^{(i)}) = \frac{1}{1+\epsilon_i(a^{(i)})} = \frac{v_i(a^{(i)})}{\beta_i(a^{(i)})}$. Hence we have :

$$\forall i \in N : \quad \varphi_i = \textstyle\sum_{a^{(i)} \in A^{(i)}} \frac{v_i(a^{(i)})}{\beta_i(a^{(i)})} z_{i,a^{(i)}} \tag{12}$$

The total payoff must be at least γ, therefore :

$$\sum_{i \in N} \sum_{a^{(i)} \in A^{(i)}} v_i(a^{(i)}) z_{i,a^{(i)}} \geq \gamma \tag{13}$$

Finally, our MIP adressing Restrict-φ(γ) is :

$$\begin{array}{ll} \text{maximize} & \min_{i \in N}\{\varphi_i\} \\ \text{subject to} & (5) \wedge (6) \wedge (7) \wedge (8) \wedge (11) \wedge (12) \wedge (13) \end{array}$$

If we assume $\frac{v_i(a^{(i)})}{\beta_i(a^{(i)})} \geq 0$ (without loss of generality), then the local-action-profile variables are consistent, and the MIP is valid. We therefore obtain MIP formulations to solve the two basic sub-problems Restrict-v(α) and Restrict-φ(γ).

Now, the greedy construction of an e-covering starts with the initial call q_0 =Restrict-v(0). Then we compute the following alternated sequences for $n \geq 1$:

$$p_n = \text{Restrict-}\varphi(q_{n-1}/(1+e_v)) \tag{14}$$

$$q_{n+1} = \text{Restrict-v}((1+e_\varphi)\varphi(p_n)) \tag{15}$$

We let n increase until the feasible domain of Restrict becomes empty. The resulting set $\{p_1, \ldots, p_K\}$ provides an e-covering set of minimal cardinality. This procedure makes only $2K$ calls to Restrict where K is the size of the output. Further details on this greedy approach and its optimality for general bi-objective problems are given in [6].

We have performed numerical tests to check the practical feasibility of this approach. The results can be seen on Figure 6 that provides e-coverings for instances of size 7×7 and $e \in \{(0.05, 0.25), (0.015, 0.075), (0.002, 0.01)\}$.

As expected, we observe that points belonging to the e-covering set spread throughout the approximate Pareto front with a density of points that increases as the norm of $e = (e_\varphi, e_v)$ diminishes. Note that the supervisor can control independently the density of points on the v-axis and on the φ-axis by choosing $e_v \neq e_\varphi$. Remark also that some of these points do not belong to the boundary of the convex-hull of the e-covering set. Such tradeoffs are known as *unsupported solutions* and cannot be obtained by optimization of a linear combination of the two criteria. The numerical tests confirm the adequacy of the e-covering concept to provide representative approximations of the Pareto ensuring some diversity

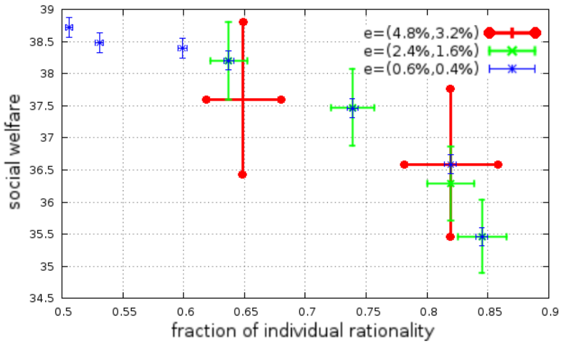

Figure 6: e-Covering of Pareto-optimal tradeoffs

(inherent to the covering property) while avoiding any redundancy (due to minimality). Such coverings provide the supervisor with a collection of typical action-profiles achieving various efficient tradeoffs between stability and social utility maximization. The final choice within the e-covering can then easily be made depending on the aspirations in terms of stability and social welfare.

6. CONCLUSION

In this paper we investigated stability and efficiency issues in GGs, without confining ourselves to the tractable case of bounded treewidth. This is made possible through the use of a variety of algorithmic solutions. Previous works showed the influence of the network topology on the asymptotic probability of PNE, exhibiting many cases where this probability tends to 0 with n. What our experimental results first exhibit is that there is an interplay between the topology and the action-set size: for some topologies only, the action-set makes a crucial difference in terms of convergence to 0 for the PNE existence. Thus, we observed a $(1-\frac{1}{e})$ probability of PNE existence in square-grids with 4 actions. To the best of our knowledge, this is so far the topology with the fastest diameter's divergence ($\sqrt{n}$) observed not to tend to 0. Our second main finding is that, despite the low probability of PNE existence, approximations of equilibria are in most cases close to individual stability. The convergences of k-Nash to particular values calls for a deeper understanding of how graph topologies influence the minimal number of unsatisfied agents. Eg., it is an open question whether there are any topology where the k-Nash over random-payoff GG may increase. While k-Nash does not rely on the payoff random uniform $[0, 1]$ distribution, ϵ-Nash does, so it would also be of interest to study non-uniform distributions.

Finally, we have seen that these approximations of PNE provide many solutions. This allows the computation of tradeoffs between almost individually stable solutions, and efficient (socially optimal) one. Knowing how the price of stability would evolve as a function of the ϵ would certainly be useful for the system designer. Our last section shows how this can be performed in practice without returning a set of undominated trade-off solutions of prohibitive size. We want to emphasize that the MIP formulations we used for the tradeoffs determinations can be generalized to other criteria. Eg., PNE can be ϵ-approximated in an additive way, or the overall utility can be replaced by a fair criterion.

7. REFERENCES

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *FOCS*, pages 295–304, 2004.

[2] U. Bertele and F. Brioschi. *Nonserial dynamic programming*. Academic Press, Inc., 1972.

[3] A. C. Chapman, A. Farinelli, E. M. de Cote, A. Rogers, and N. R. Jennings. A distributed algorithm for optimising over pure strategy Nash equilibria. In *AAAI*, 2010.

[4] C. Daskalakis, A. Dimakis, and E. Mossel. Connectivity and equilibrium in random games. *The Annals of Applied Probability*, 21(3):987–1016, 2011.

[5] C. Daskalakis and C. Papadimitriou. Computing pure Nash equilibria in graphical games via markov random fields. In *ACM-EC*, pages 91–99, 2006.

[6] I. Diakonikolas and M. Yannakakis. Small approximate Pareto sets for biobjective shortest paths and other problems. *SIAM Journal on Computing*, 39(4):1340–1371, 2009.

[7] B. N. Dilkina, C. P. Gomes, and A. Sabharwal. The impact of network topology on pure Nash equilibria in graphical games. In *AAAI*, pages 42–49, 2007.

[8] N. Een and N. Sorensson. An extensible SAT-solver. In E. Giunchiglia and A. Tacchella, editors, *SAT*, volume 2919 of *LNCS*, pages 502–518. Springer, 2003.

[9] Y. Gao. Treewidth of Erdős-rényi random graphs, random intersection graphs, and scale-free random graphs. *Discrete Applied Mathematics*, 2011.

[10] G. Gottlob, G. Greco, and F. Scarcello. Pure Nash equilibria: Hard and easy games. In *TARK-03*, pages 215–230. ACM, 2003.

[11] I. Gurobi Optimization. Gurobi optimizer reference manual, 2012.

[12] A. Jiang and M. Safari. Pure Nash equilibria: complete characterization of hard and easy graphical games. In *AAMAS*, pages 199–206, 2010.

[13] M. Kearns. Graphical games. In *Algorithmic Game Theory*, chapter 7. Cambridge, UK: Cambridge University Press, 2007.

[14] M. J. Kearns, M. L. Littman, and S. P. Singh. Graphical models for game theory. In *UAI*, 2001.

[15] J. M. Mooij. libDAI: A free and open source C++ library for discrete approximate inference in graphical models. *JMLR*, 11:2169–2173, Aug. 2010.

[16] C. Papadimitriou and M. Yannakakis. On the approximability of trade-offs and optimal access of web sources. In *FOCS*, pages 86–92, 2000.

[17] Y. Sharma and D. P. Williamson. Stackelberg thresholds in network routing games or the value of altruism. *GEB*, 67(1):174–190, 2009.

[18] Y. Shoham and K. Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.

[19] V. Soni, S. Singh, and M. Wellman. Constraint satisfaction algorithms for graphical games. In *AAMAS*, page 67. ACM, 2007.

[20] D. Vickrey and D. Koller. Multi-agent algorithms for solving graphical games. In *AAAI*, 2002.