10171 Abstracts Collection
Equilibrium Computation
— Dagstuhl Seminar —

Edith Elkind[1], Nimrod Megiddo[2], Peter Bro Miltersen[3], Bernhard von Stengel[4]
and Vijay Vazirani[5]

[1] Nanyang Technological University, SG
`eelkind@ntu.edu.sg`
[2] IBM Almaden Center – San José, US
`megiddo@theory.stanford.edu`
[3] University of Aarhus, DK
`bromille@cs.au.dk`
[4] London School of Economics, UK
`stengel@maths.lse.ac.uk`
[5] Georgia Institute of Technology, US
`vazirani@cc.gatech.edu`

Abstract. From April 25 to April 30, 2010, the Dagstuhl Seminar 10171
"Equilibrium Computation" was held in Schloss Dagstuhl – Leibniz Cen-
ter for Informatics. During the seminar, several participants presented
their current research, and ongoing work and open problems were dis-
cussed. Abstracts of the presentations given during the seminar as well
as abstracts of seminar results and ideas are put together in this paper.
The first section describes the seminar topics and goals in general. Links
to extended abstracts or full papers are provided, if available.

Keywords. Equilibrium computation, algorithmic game theory

10171 Executive Summary – Equilibrium Computation

The focus of this seminar was the algorithmic problem of computing equilibria
in games and market models, viewed from both the theoretical and practical
perspective. The equilibrium computation problem is one of the central topics
in the rapidly expanding field of algorithmic game theory.

The seminar was a follow-up to Seminar 07471, on the same topic, but with
three new organizers, and with a focus on some of the aspects of this prob-
lem that received relatively little attention in Seminar 07471. One of the major
themes of this seminar was *dynamics*, i.e., exploring the agents' behavior (at
both individual and collective level) that leads to the discovery of equilibria,
and, more generally, adaptive changes in the collective behavior. Discussed were
the classic game-theoretic approaches to this topic (organizer: von Stengel) as
well as more recent computational and simulation-based techniques, as stud-
ied by the multi-agent community (organizer: Elkind). Another key emphasis

was on algorithms and complexity results for *market equilibria* and their applications to Nash Bargaining Games (organizer: Vazirani). We also compared these approaches with computational and geometric aspects of the central Linear Complementarity Problem (LCP) in mathematical programming (organizer: Megiddo). Finally, since the last seminar there was significant progress understanding the complexity of important algorithms, such as strategy iteration, for solving two-player zero-sum games of infinite duration. Progress in this area has strong connections to mathematical programming and was discussed and extended (organizer: Miltersen).

The following abstracts indicate the diversity of topics and their rich interconnections that were explored during a very successful seminar.

Viewing market price discovery as an algorithmic process

Richard Cole (New York University, US)

A basic tenet of well-functioning markets is that they converge toward equilibrium prices. This talk seeks to explain why this might be possible.

More specifically, we introduce the setting of Ongoing Markets (by contrast with the classic Fisher and Exchange markets). An Ongoing Market allows trade at non-equilibrium prices, and, as its name suggests, continues over time. Also, we specify and analyze a (tatonnement-style) price update rule with the following characteristics:
1. There is a separate instance of the (price update) procedure for each good.
2. The procedure is distributed: (i) the instances are independent, and (ii) each instance uses limited "local" information.
3. It is simple.
4. It is asynchronous: price updates do not have to be simultaneous.
   Furthermore, for appropriate markets the rule enables:
5. Fast convergence.
6. Robustness in the face of (somewhat) inaccurate data.
   This talk is based on joint works with Lisa Fleischer and Ashish Rastogi.

A generalization of the min-max theorem to multiplayer games

Constantinos Daskalakis (MIT – Cambridge, US)

According to Aumann, zero-sum games are "one of the few areas in game theory, and indeed in the social sciences, where a fairly sharp, unique prediction is made". We provide a generalization of the minmax theorem to multi-player games. The games we consider are zero-sum polymatrix—that is, every pair of players plays a (potentially different) two-player game, and every outcome of the global interaction has zero sum of all players' payoffs. Our generalization of

the minmax theorem to this setting implies convexity of equilibria, polynomial-time tractability, and convergence of no-regret learning algorithms to equilibria. We show that our class of games is essentially the broadest class of multiplayer zero-sum games to which we can hope to push tractability results.

And what about extending our results beyond zero-sum games? Previous work has established that computing exact Nash equilibria is computationally intractable, and research on approximation algorithms has made no progress beyond finite values of approximation. Nevertheless, inapproximability results have been evading current techniques. We provide the first inapproximability result for Nash equilibria in two-player games, for constant values of relative approximation.

Keywords: Equilibrium Computation

Competitive Equilibrium at Advertising Marketplaces

Xiaotie Deng (City University – Hong Kong, CN)

We apply the competitive equilibrium formulation to study an advertising market problem where each position has an associated parameter called circulation, quantifying the worth of the position, e.g., the number of viewers in TV advertising, or the average amount of clicks for each 1000 impressions in keywords advertising. Each advertiser has a value constraint which refers to his maximum amount for one unit circulation, and a total budget constraint for all circulations in a position. In an output of the market, a matching between advertisers and positions is established and a price-per-circulation of each position is charged.

The problem is known not to guarantee an integral equilibrium. Nevertheless of the non-existence, we develop a strongly polynomial time algorithm to determine whether there is an equilibrium, and to compute a minimum one, where the total payment of advertisers is the minimum. Our algorithm can be extended to find an equilibrium with an opposite objective that maximizes the total profit of the publisher. Further, we establish the existence of an equilibrium, in particular when the budgets of advertisers are distinct, a mild condition that can, arguably, be reenforced. Finally, we consider a natural mechanism based on the minimum equilibrium and prove it is incentive compatible.

Our results open up the possibility of a new solution for pricing an advertising market by its general equilibrium, and shed new light on possible algorithmic solutions for the emerging ad exchanges.

Keywords: Equilibrium Computation

Joint work of: Chen, Ning; Deng, Xiaotie

A new convex program for Fisher markets and convergence of proportional response dynamics

Nikhil R. Devanur (Microsoft Research – Redmond, US)

We give a new convex program that captures the equilibrium of linear Fisher markets.

This is an alternate to the well known Eisenberg-Gale convex program that has been known and studied for a long time (since 1959).

This new program naturally extends to the case of spending constraint utilities, thus resolving the open problem of whether there exists a convex program for such utilities. The new program also helps us demystify the convergence properties of "proportional response" dynamics: a simple, distributed way to converge to market equilibria.

We show that the proportional response dynamics is equivalent to gradient descent on the new convex program, with KL-divergence instead of Euclidean distance. In addition to providing insight, this allows us to give improved bounds on the convergence rate.

Manuscript available from

Full Paper:
  http://research.microsoft.com/en-us/um/people/nikdev/

Equilibria of Plurality Voting with Abstentions

Edith Elkind (Nanyang TU – Singapore, SG)

In the traditional voting manipulation literature, it is assumed that a group of manipulators jointly misrepresent their preferences to get a certain candidate elected, while the remaining voters are truthful.

In this paper, we depart from this assumption, and consider the setting where all voters are strategic.

In this case, the election can be viewed as a game, and the election outcomes correspond to Nash equilibria of this game. We use this framework to analyze two variants of Plurality voting, namely, simultaneous voting, where all voters submit their ballots at the same time, and sequential voting, where the voters express their preferences one by one. For simultaneous voting, we characterize the preference profiles that admit a pure Nash equilibrium, but show that it is computationally hard to check if a given profile fits our criterion.

For sequential voting, we provide a complete analysis of the setting with two candidates, and show that for three or more candidates the equilibria of sequential voting may behave in a counterintuitive manner.

Keywords: Equilibrium Computation

Joint work of: Desmedt, Yvo; Elkind, Edith

Adding Recursion to Markov Chains, Markov Decision Processes, and Stochastic Games: Algorithms and Complexity

Kousha Etessami (University of Edinburgh, GB)

I will decribe a family of finitely presented, but infinite-state, stochastic models and stochastic games that arise by adding a natural recursion feature to Markov Chains, Markov Decision Processes, and Stochastic Games.

These recursive models subsume a number of classic and heavily studied stochastic processes, including (multi-type) branching processes, (quasi-)-birth-death processes, stochastic context-free grammars, and many others. They also provide a natural abstract model of probabilistic procedural programs with recursion.

The algorithmic theory and computational complexity of analyzing these models has turned out to be very rich, with connections to a number of areas of research. In particular, a key role is played in their analysis by fixed point computation problems for algebraically defined functions over basis +, *, max, min, and there are connections between some of these problems and long standing open problems in numerical computation, as well as connections to the computation and approximation of Nash Equilibria.

I will survey highlights from our work on these recursive stochastic models and games. There remain many open questions about the complexity of basic analysis problems for them. I will highlight a few key open problems.

Keywords: Equilibrium Computation

Joint work of: Etessami, Kousha; Yannakakis, Mihalis

Exponential Lower Bounds for Strategy Iteration

Oliver Friedmann (LMU München, DE)

We present a new exponential lower bound for the two most popular deterministic variants of the strategy improvement algorithm for solving parity, mean payoff, discounted payoff and simple stochastic games. The first variant improves every node in each step maximizing the current valuation locally, whereas the second variant computes the optimal improvement in each step. We outline families of games on which both variants require exponentially many strategy iterations. Finally, we briefly describe how to adapt the idea to get a concrete lower bound for the random facet policy.

Keywords: Policy iteration, parity games, lower bounds

Stochastic games with signals: decidable and undecidable problems

Hugo Gimbert (Université Bordeaux, FR)

We consider two-player stochastic games with signals.

These games are similar to the original stochastic games introduced by Shapley except for the signals.

These signals are used to model imperfect monitoring: the players cannot observe directly the state of the game, instead they base their decisions on the sequence of signals they receive during the play.

We present several decidability and undecidability results for these games equipped with reachability conditions: for example values are not computable but it is decidable whether a player has a strategy to win almost-surely.

Keywords: Equilibrium Computation

A perfect price discrimination market, its welfare theorems, and an efficient algorithm for computing its equilibria

Gagan Goel (Georgia Institute of Technology, US)

We present a simple market model which allows for perfect price discrimination and we give a notion of equilibrium for it. We also give a convex program, a generalization of the classic Eisenberg-Gale convex program, that captures all equilibria for this market. This program gives us surprisingly simple proofs of both welfare theorems for this market. Finally, we give a combinatorial, polynomial time algorithm for finding an equilibrium and a characterization of the set of all equilibria.

Joint work of: Goel, Gagan; Vazirani, Vijay

Nucleolus Computation in Compact Coalitional Games

Gianluigi Greco (University of Calabria, IT)

A coalitional game is played by a set of players that can form coalitions in order to guarantee themselves some advantage. Each coalition is assigned a certain worth that players in it can obtain if collaborating with each other, and the outcome of the game is a vector of payoffs that is meant to specify the distribution of the worth granted to each player in the game. A fundamental problem for coalitional games is to define the most desirable outcomes in terms of appropriate notions of fair worth distributions. The talk focuses on characterizing the

computational complexity of the nucleolus, which is a very prominent solution concept of fairness in coalitional games. In particular, we consider games that are specified in succinct form (i.e., the worth function is given as polynomial-time oracle) and we show that the nucleolus can be computed in $F\Delta_2^P$ in this setting. Thus, checking whether a given imputation is the nucleolus of a succinctly specified game is feasible in Δ_2^P. In fact, we show that the result is tight, since Δ_2^P-hardness holds even on the class of graph games.

Keywords: Coalitional games, Computational Complexity

Full Paper:
  http://ijcai.org/papers09/Papers/IJCAI09-035.pdf

A New Result on the Complex of Maximal Lattice Free Bodies Arising from a matrix of size $(n+1) \times n$

Bjarke Hammersholt Roune (Aarhus University, DK)

We consider the collection of Maximal Lattice Free Bodies arising from a well behaved matrix of size $(n+1)$ by n. When $n=2$, the bodies consist of a pair of triangles, whose union is a parallelogram of unit area – and all of their lattice translates. When $n=3$, the bodies are tetrahedra, whose vertices are integral and which contain no other lattice points in their interior or on their boundary. A theorem from 1985, with an impossible opaque argument, states that all of these tetrahedra have a lattice width of 0 or 1 in the SAME lattice direction.

We present a simple argument based on the new result that the collection of MLFBs contains a pair of triangles $(0, a, a+b)$, $(0, b, a+b)$.

Joint work of: Scarf, Herbert

Hardness of approximating minmax in 3-player games

Kristoffer Arnsfelt Hansen (Aarhus University, DK)

For computing the minmax value for a given player in 3-player normal form games we provide the following results: For any constant $\epsilon > 0$, the minmax value can be approximated to an additive error $\epsilon > 0$ in quasipolynomial time. It is known that obtaining an additive error $1/n^\epsilon$ for any $\epsilon > 0$ is NP-hard. We give a reduction from the k-clique problem, which additionally implies W[1] hardness if the problem is paramerized by the number k of strategies for the given player. Finally, assuming hardness of the so-called planted clique problem, we show that it is not possible to approximate the minmax value to any constant additive error $\epsilon > 0$ in polynomial time.

Joint work of: Hansen, Kristoffer Arnsfelt; Hansen, Thomas Dueholm; Miltersen, Peter Bro; Sørensen, Troels Bjerre; Verbin, Elad

A subexponential lower bound for the Random Facet algorithm for Parity Games

Thomas Dueholm Hansen (Aarhus University, DK)

Parity Games form an intriguing family of infinite duration games whose solution is equivalent to the solution of important problems in automatic verification and automata theory. They also form a very natural subclass of Deterministic Mean Payoff Games, which in turn is a very natural subclass of turn-based Stochastic Mean Payoff Games. It is a major open problem whether these game families can be solved in polynomial time.

The currently fastest algorithms for the solution of all these games are adaptations of the randomized algorithms of Kalai and of Matousek, Sharir and Welzl for LP-type problems, an abstract generalization of linear programming. We refer to the algorithm of Matousek, Sharir and Welzl as the Random Facet algorithm. The expected running time of these algorithms is subexponential in the size of the game, i.e., $2^{O(\sqrt{n \log n})}$, where n is the number of vertices in the game. Matousek constructed a family of abstract optimization problems such that the expected running time of the Random Facet algorithm, when run on a random instance from this family, is close to the subexponential upper bound given above. This shows that in the abstract setting, the $2^{O(\sqrt{n \log n})}$ upper bound on the complexity of the Random Facet algorithm is essentially tight.

It is not known, however, whether the abstract optimization problems constructed by Matousek correspond to games of any of the families mentioned above. (It is, in fact, unlikely that they do.) There was some hope, therefore, that the Random Facet algorithm, when applied to, say, parity games, may run in polynomial time. We show, unfortunately, that this is not the case by constructing explicit parity games on which the expected running time of the Random Facet algorithm is close to the subexponential upper bound.

Our parity games are obtained by modifying the parity games used recently by Friedmann to show that the standard deterministic policy iteration algorithm for parity games runs in exponential time. They, in a sense, implement a randomized counter. They are also the first explicit LP-type problems on which the Random Facet algorithm is not polynomial.

Joint work of: Friedmann, Oliver; Hansen, Thomas Dueholm; Zwick, Uri

Dynamics and Equilibria

Sergiu Hart (The Hebrew University at Jerusalem, IL)

It is a fact that in the existing literature there are no general natural dynamics leading to Nash equilibria. The talk provides an overview of research that sheds light on this and related issues.

We provide an overview of recent work on dynamical systems in multi-player environments. On the one hand, the natural informational restriction that each

participant does not know the payoff functions of the other participants – "uncoupledness" – severely limits the possibilities to converge to Nash equilibria. On the other hand, there are simple adaptive heuristics – such as "regret matching" – that lead in the long run to correlated equilibria, a concept that embodies full rationality. Connections to behavioral economics, neurobiological studies, and engineering, are also mentioned.

Full Paper:
  http://www.ma.huji.ac.il/hart/abs/dyn-p.html

Solving Simple Stochastic Tail Games

Florian Horn (University Paris-Diderot, FR)

Stochastic games are a natural model for open reactive processes: one player represents the controller and his opponent represents a hostile environnment. Teh evolution of the system depends on the decisions of the players, supplemented by random transitions. There are two main algorithmic problems on such games: computing the values (quantitative analysis) and deciding whether a player can win with probability 1 (qualitative analysis). In this paper we reduce the quantitative analysis of simple stochastic tail games to the qualitative analysis: we provide an algorithm for computing values which uses qualitative analysis as a sub-procedure. The correctness proof of this algorithm reveals several nice properties of perfect-information stochastic tail games, in particular the existence of optimal strategies. We apply these results to games whose winning conditions are boolean combinations of mean-payoff and Büchi conditions.

Keywords: Equilibrium Computation

Joint work of: Gimbert, Hugo; Horn, Florian

Lemke's algorithm for discounted games

Marcin Jurdzinski (University of Warwick, GB)

The performance of a pivoting algorithm due to Lemke is considered on linear complementarity problems (LCPs) that arise from discounted games. The algorithm has not been previously studied in the context of infinite games, and it offers an alternative to the classical strategy-improvement algorithms. The algorithm is described purely in terms of discounted games, thus bypassing the reduction from the games to LCPs, and hence facilitating a better understanding of the algorithm when applied to games. A family of discounted games is given on which the algorithm runs in exponential time, indicating that in the worst case it performs no better for discounted games than it does for general P-matrix LCPs.

Keywords: Equilibrium Computation

Computational Mechanism Analysis: Leveraging Equilibrium Computation to Understand Real-World Mechanisms

Kevin Leyton-Brown (University of British Columbia – Vancouver, CA)

Many mechanisms become important "in the wild" despite a lack of theoretical arguments in their favor, chiefly because their complexity precludes analysis by existing techniques. This talk advocates the creation of computational tools for describing and empirically analyzing the equilibria of such mechanisms. This agenda raises a host of new theoretical problems: identifying representations that compactly represent interesting classes of games; determining encodings of mechanisms of interest into these representations; and deriving efficient procedures for computing solution concepts of interest given these encodings. Recently, we have begun to take steps to make this agenda concrete. Specifically, I will describe "Action Graph Games", a compact encoding for perfect-information games in which agents' payoffs exhibit context-free independencies. I will show how this representation can be used to encode sponsored search mechanisms, and will characterize resulting equilibrium behavior both in terms of social welfare and revenue. I will also outline some recent extensions of our representation to encompass Bayesian games.

Keywords: Equilibrium Computation

Coalitional Bargaining; How Complex Is It?

Andrew McLennan (The University of Queensland – Brisbane, AU)

We consider a model of coalitional bargaining in which, in each period, a proposer is selected randomly, the proposer suggests a division of a pie of size one, and there is an election in which each agents votes for or against the proposal. If the set of agents voting for the proposal is a winning coalition, the proposal is implemented, and otherwise the process is repeated in the next period with discounted payoffs. We show that there is a unique vector of expected equilibrium payoffs generated by every stationary equilibrium. The proof uses index theory, and does not imply an algorithm for computing the vector. Verification of the vector is a linear program, so the problem of computing the vector is NP-like (possibly difficult to find, but easy to verify) but the equations are nonlinear, so the components of the vector need not be rational.

Keywords: Coalitional bargaining, FIXP, PPAD

Joint work of: McLennan, Andrew; Hulya Eraslan

Finding Gale Strings

Julian Merschen (London School of Economics, GB)

The problem 2-NASH of finding a Nash equilibrium of a bimatrix game belongs to the complexity class PPAD. This class comprises computational problems that are known to have a solution by means of a path-following argument. For bimatrix games, this argument is provided by the Lemke-Howson algorithm. It has been shown that this algorithm is worst-case exponential with the help of dual cyclic polytopes, where the algorithm can be expressed combinatorially via labeled bitstrings defined by the 'Gale evenness condition' that characterize the vertices of these polytopes. We define the combinatorial problem ANOTHER COMPLETELY LABELED GALE STRING whose solutions define the Nash equilibria of games defined by cyclic polytopes, including games where the Lemke-Howson algorithm takes exponential time. If this problem was PPAD-complete, this would imply that 2-NASH is PPAD-complete, in a much simpler way than the currently known proofs, including the original proof by Chen and Deng [3]. However, we show that ANOTHER COMPLETELY LABELED GALE STRING is solvable in polynomial time by a simple reduction to PERFECT MATCHING in graphs, making it unlikely to be PPAD-complete. Although this result is negative, we hope that it stimulates research into combinatorially defined problems that are PPAD-complete and imply this property for 2-NASH.

Keywords: Nash Equilibria, Lemke-Howson, PPAD, Perfect Matching, Oiks, Gale Strings

Joint work of: Casetti, Marta; Merschen, Julian; von Stengel, Bernhard

Full Paper:
 http://www.maths.lse.ac.uk/Personal/stengel/TEXTE/ge-ISCO.pdf

The complexity of solving reachability games using value and strategy iteration

Peter Bro Miltersen (Aarhus University, DK)

Concurrent reachability games is a class of games heavily studied by the computer science community, in particular by the formal methods community. Two standard algorithms for approximately solving two-player zero-sum concurrent reachability games are value iteration and strategy iteration. A rigorous complexity analysis of these algorithms has been an open problem until now. We prove a lower bound of $2^{m^{\Omega(N)}}$ on the worst case number of iterations needed for both of these algorithms to provide non-trivial approximations to the value of a game with N non-terminal positions and m actions for each player in each

position. In particular, both algorithms have at least doubly-exponential complexity. Also, even when the game given as input has only one non-terminal position, we prove an exponential lower bound on their time complexities. The instances establishing the lower bound may be regarded as natural rather than pathological and our proofs of the lower bounds proceed by arguing about these instances as games, using game-theoretic concepts and tools. We also give not-quite-matching $2^{2^{O(mN)}}$ upper bounds on the number of iterations sufficient for getting adequate approximate values, by each of these algorithms. In particular, both algorithms are also of at most doubly-exponential complexity.

Keywords: Equilibrium Computation

Joint work of: Hansen, Kristoffer Arnsfelt; Ibsen-Jensen, Rasmus; Miltersen, Peter Bro

The P-matrix linear complementarity problem: survey of complexity results

Walter Morris (George Mason University – Fairfax, US)

The P-matrix linear problem is that of finding, given a square matrix M and vector q, either a solution to the linear complementarity problem with input M and q or a nonpositive principal minor of the matrix M. This problem is in PPAD, and Megiddo showed that NP-hardness of the problem would imply NP = Co-NP. A large variety of complexity classes is associated to different specializations of this problem. Some instances of this problem can be solved by considering a related linear program, as was observed by Mangasarian and Cottle and Pang in the 1970s. We contrast those instances for which this and related approaches work with the instances for which they do not.

Improved Algorithms for Computing Fisher's Market Clearing Prices

James B. Orlin (MIT – Cambridge, US)

We give the first strongly polynomial time algorithm for computing an equilibrium for the linear utilities case of Fisher's market model.
   We consider a problem with a set B of buyers and a set G of divisible goods.
   Each buyer i starts with an initial integral allocation e_i of money. The integral utility for buyer i of good j is U_{ij}. We first develop a weakly polynomial time algorithm that runs in $O(n^4 \log U_{max} + n^3 e_{max})$ time, where $n = |B| + |G|$. We further modify the algorithm so that it runs in $O(n^4 \log n)$ time. These algorithms improve upon the
   previous best running time of $O(n^8 \log U_{max} + n^7 \log e_{max})$, due to Devanur et al.

Computing equilibria: the plot thickens

Christos H. Papadimitriou (University of California – Berkeley, US)

We show new complexity results regarding Nash equilibria in the face of risk, equilibrium selection à la Harsanyi-Selten, and equilibria in economies with con-concave production.

Congestion-Averse Games

Maria Polukarov (University of Southampton, GB)

Congestion games – in which players strategically choose from a set of "resources" and derive utilities that depend on the congestion on each resource – are important in a wide range of applications. However, to date, such games have been constrained to use utility functions that are linear sums with respect to resources. To remove this restriction, we consider a generalisation to the case where a player's payoff can be given by any real-valued function over the set of possible congestion vectors. Under weak assumptions on the structure of player strategy spaces, we constructively prove the existence of a pure strategy equilibrium for the very wide class of these generalised games in which player utility functions are congestion-averse – i.e., monotonic, submodular and independent of irrelevant alternatives. Although, as we show, these games do not admit a generalised ordinal potential function (and hence – the finite improvement property), any such game does possess a Nash equilibrium in pure strategies. A polynomial time algorithm for computing such an equilibrium is presented.

Moreover, we show that if a player utility function does not satisfy any one of the congestion-averse conditions, then a pure strategy equilibrium need not exist, and in fact the determination of whether or not a game has a pure strategy Nash equilibrium is NP-complete. We further prove analogous results for our assumed properties of player strategy spaces, thus showing that current assumptions on strategy and utility structures in this model cannot be relaxed anymore.

Computing stable outcomes in hedonic games

Rahul Savani (University of Liverpool, GB)

We study the computational complexity of finding stable outcomes in additively-separable symmetric hedonic games. These coalition formation games are specified by an undirected edge-weighted graph: nodes are players, an outcome of the game is a partition of the nodes into coalitions, and the utility of a node is the sum of incident edge weights in the same coalition. We consider several natural stability requirements defined in the economics literature. For all of them the existence of a stable outcome is guaranteed by a potential function argument, and local improvements will converge to a stable outcome. The different stability requirements correspond to different local search neighbourhoods. For different neighbourhood structures, our findings comprise positive results in the form of polynomial-time algorithms, and negative (PLS-completeness) results.

Keywords: Equilibrium Computation

Joint work of: Gairing, Martin; Savani, Rahul

Computing Pure Nash and Strong Equilibria in Bottleneck Congestion Games

Alexander Skopalik (RWTH Aachen, DE)

Bottleneck congestion games properly model the properties of many real-world network routing applications. They are known to possess strong equilibria – a strengthening of Nash equilibrium to resilience against coalitional deviations. In this paper, we study the computational complexity of pure Nash and strong equilibria in these games. We provide a generic centralized algorithm to compute strong equilibria in bottleneck congestion games. It has polynomial running time for many interesting classes of games such as, e.g., matroid or single-commodity games.

In addition, we examine the more demanding goal to reach equilibria in polynomial time using natural improvement dynamics. Using unilateral improvement dynamics in matroid games pure Nash equilibria can be reached efficiently. In contrast, computing even a single coalitional improvement move in matroid and single-commodity games is strongly NP-hard. In addition, we establish a variety of hardness results and lower bounds regarding the duration of unilateral and coalitional improvement dynamics. They are extended to hold even for convergence to approximate equilibria.

Joint work of: Harks, Tobias; Hoefer, Martin; Klimm, Max; Skopalik, Alexander

The computational complexity of trembling hand perfection and other equilibrium refinements

Troels Bjerre Sorensen (University of Warwick, GB)

The king of refinements of Nash equilibrium is trembling hand perfection. We show that it is NP-hard and SQRT-SUM-hard to decide if a given pure strategy Nash equilibrium of a given three-player game in strategic form with integer payoffs is trembling hand perfect. Analogous results are shown for a number of other solution concepts, including proper equilibrium, (the strategy part of) sequential equilbrium, quasi-perfect equilibrium and CURB. The proofs all use a reduction from the problem of comparing the minmax value of a three-player game in strategic form to a given rational number. This problem was previously shown to be NP-hard by Borgs et al., while a SQRT-SUM hardness result is given in this paper. The latter proof yields bounds on the algebraic degree of the minmax value of a three-player game that may be of independent interest.

We remark that the question answered by this paper was raised during the previous Dagstuhl seminar on Equilibrium computation (seminar 07471) .

Keywords: Equilibrium Computation

Joint work of: Hansen, Kristoffer Arnsfelt; Miltersen, Peter Bro; Sorensen, Troels Bjerre

An algorithmic proof of Afriat's theorem on consumer behavior and related questions

Michael J. Todd (Cornell University, US)

We give an algorithmic proof of a theorem of Afriat characterizing when a set of demand data arises from utility maximization, and discuss related questions. Afriat's theorem shows that some obviously necessary conditions are also sufficient. We give an inductive proof that leads to an algorithm that either produces a possible utility function or shows that none exists in $O((m+n)n2)$ operations, where there are m goods and n demand vector-price vector pairs.

Keywords: Consumer behavior; demand rationalization; utility maximixation

Joint work of: Todd, Michael; Fostel, Ana; Scarf, Herb

Full Paper:
  http://www.springerlink.com/content/b2da6l6dvjljxrl0/?p=62c6c53692eb432f9d28ddc112209424

Proportional Response as Iterated Cobb-Douglas

Michael J. Todd (Cornell University, US)

We show that the proportional response algorithm for computing an economic equilibrium in a Fisher market model can be interpreted as iteratively approximating the economy by one with Cobb-Douglas utilities, for which a closed-form equilibrium can be obtained.

We also extend the method to allow elasticities of substitution at most one.

Keywords: Computing equilibria, Fisher market, proportional response algorithm, Cobb-Douglas utilities

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2010/2671

Market Equilibrium: the Quest for the "Right" Model

Vijay V. Vazirani (Georgia Institute of Technology, US)

The search for efficient algorithms for computing market equilibria had not only an academic goal, i.e., providing an algorithmic ratification of Adam Smith's "invisible hand of the market," but was also motivated by potential applications to the plethora of new and highly lucrative markets that have emerged on the Internet.

However, this program was recently dealt a serious blow, with results showing that the problem of computing an equilibrium under even additively separable, piecewise-linear, concave utilities (plc utilities) is PPAD-complete for both Arrow-Debreu and Fisher market models. Assuming P $\neq$ PPAD, this effectively rules out the existence of efficient algorithms for almost all general and interesting classes of "traditional" market models.

On the other hand, markets in the West, based on Adam Smith's free market principle, seem to do a good job of finding prices that maintain parity between supply and demand. This has prompted the question of whether we have failed to capture some essential elements of real markets in our models, and what is the "right" model which is not only realistic and admits equilibria but is also amenable to efficient computation of equilibria.

We will give evidence that there is plenty of scope of doing creative work on this question.

Full Paper:
 http://www.cc.gatech.edu/~vazirani/Dagstuhl-2010.ppt

On the Complexity of Nash Equilibria and Other Fixed Points

Mihalis Yannakakis (Columbia University, US)

We reexamine what it means to compute Nash equilibria and, more generally, what it means to compute a fixed point of a given Brouwer function, and we investigate the complexity of the associated problems. Specifically, we study the complexity of the following problem: given a finite game, Γ, with 3 or more players, and given $\epsilon > 0$, compute an approximation within ϵ of some (actual) Nash equilibrium. We show that approximation of an actual Nash Equilibrium, even to within any non-trivial constant additive factor $\epsilon < 1/2$ in just one desired coordinate, is at least as hard as the long standing square-root sum problem, as well as a more general arithmetic circuit decision problem that characterizes P-time in a unit-cost model of computation with arbitrary precision rational arithmetic; thus placing the approximation problem in P, or even NP, would resolve major open problems in the complexity of numerical computation.

We show similar results for market equilibria: it is hard to estimate with any nontrivial accuracy the equilibrium prices in an exchange economy with a unique equilibrium, where the economy is given by explicit algebraic formulas for the excess demand functions.

We define a class, FIXP, which captures search problems that can be cast as fixed point computation problems for functions represented by algebraic circuits (straight line programs) over basis $\{+, *, -, /, max, min\}$ with rational constants. We show that the (exact or approximate) computation of Nash equilibria for 3 or more players is complete for FIXP. The price equilibrium problem for exchange economies with algebraic demand functions is another FIXP-complete problem. We show that the piecewise linear fragment of FIXP equals PPAD.

Many other problems in game theory, economics, and probability theory, can be cast as fixed point problems for such algebraic functions. We discuss several important such problems: computing the value of Shapley's stochastic games, and the simpler games of Condon, extinction probabilities of branching processes, probabilities of stochastic context-free grammars, and termination probabilities of Recursive Markov Chains. We show that for some of them, the approximation, or even exact computation, problem can be placed in PPAD, while for others, they are at least as hard as the square-root sum and arithmetic circuit decision problems.

Homotopy methods for finding quantal response equilibria

Boyu Zhang (Universität Wien, AT)

Following Mckelvey and Palfrey, we study the quantal response equilibria for 2×2 bimatrix games. Based on homotopy methods, we find a simple rule to decide which equilibrium is selected for given payoff matrix. Generally, this rule is independent of the distribution of the random error.

Keywords: Quantal response equilbria, homotopy methods, equilibrium selection

Local Improvement/Policy Iteration Algorithms

Uri Zwick (Tel Aviv University, IL)

Policy Iteration algorithms form a simple family of algorithms that can be applied in many different settings, ranging from the relatively simple problem of finding a minimum mean weight cycle in a graph, the more challenging solution of Markov Decision Processes (MDPs), to the solution of 2-player full information stochastic games, also known as Simple Stochastic Games (SSGs).

It was recently shown by Fridmann that the worst case running time of a natural deterministic version of the policy iteration algorithm, when applied to Parity Games (PGs), is exponential. It is still open, however, whether deterministic policy iteration algorithm can solve Markov Decision Processes in polynomial time, and whether randomized policy iteration algorithms can solve Simple Stochastic Games in polynomial time.

The talk will survey what is known regarding policy iteration algorithms and mention many intriguing open problems.

Keywords: Equilibrium Computation