

07271 Abstracts Collection
Computational Social Systems and the Internet
— **Dagstuhl Seminar** —

Peter Cramton¹, Rudolf Müller², Eva Tardos³ and Moshe Tennenholtz⁴

¹ Univ. of Maryland - College Park, USA

cramton@umd.edu

² Maastricht Univ., NL

³ Cornell Univ., US

eva@cs.cornell.edu

⁴ Technion - Haifa, IL

moshet@ie.technion.ac.il

Abstract. From 01.07. to 06.07.2007, the Dagstuhl Seminar 07271 “Computational Social Systems and the Internet” was held in the International Conference and Research Center (IBFI), Schloss Dagstuhl. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. 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. Mechanism Design, Combinatorial Auctions, Social Choice Theory, Behavioral Economics, Computational Game Theory, Social Networks

07271 Summary – Computational Social Systems and the Internet

The seminar "Computational Social Systems and the Internet" facilitated a very fruitful interaction between economists and computer scientists, which intensified the understanding of the other disciplines' tool sets.

The seminar helped to pave the way to a unified theory of social systems on the Internet that takes into account both the economic and the computational issues—and their deep interaction.

Keywords: Mechanism Design, Combinatorial Auctions, Social Choice Theory, Behavioral Economics, Computational Game Theory, Social Networks

Joint work of: Cramton, Peter; Müller, Rudolf; Tardos, Eva; Tennenholtz, Moshe

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1164>

An Axiomatic Approach to Personalized Ranking Systems

Alon Altman (Technion - Haifa, IL)

Personalized ranking systems and trust systems are an essential tool for collaboration in a multi-agent environment. In these systems, trust relations between many agents are aggregated to produce a personalized trust rating of the agents. In this paper we introduce the first extensive axiomatic study of this setting, and explore a wide array of well-known and new personalized ranking systems. We adapt several axioms (basic criteria) from the literature on global ranking systems to the context of personalized ranking systems, and fully classify the set of systems that satisfy all of these axioms. We further show that all these axioms are necessary for this result.

Keywords: Ranking systems, trust, axiomatization, incentives, mechanism design, game theory

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2007/1152>

Full Paper:

http://www.technion.ac.il/~alon_a/personalized.pdf

Robust Learning Equilibrium

Itai Ashlagi (Technion - Haifa, IL)

A learning equilibrium is a profile of learning algorithms, in which a unilateral deviation from a learning algorithm is not-beneficial regardless of the state of nature.

We are especially interested in a learning equilibrium which leads to optimal surplus, in which both convergence to optimal surplus and the number of stages needed for any deviation to become non-beneficial are polynomial. In the first part of this talk, we enrich the class of settings for which such learning equilibrium exists.

We consider a resource selection game with incomplete information about the resource-cost functions. All the players know is the set of players, an upper bound on the possible costs, and that the cost functions are positive and nondecreasing. The game is played repeatedly and after every stage each player observes her cost, and the actions of all players. For such settings, we prove the existence of a learning equilibrium, yielding optimal surplus.

In the second part of the lecture we introduce the study of robust learning equilibrium - a learning equilibrium which is immune to failure of the agents to follow their algorithms for some finite time. We consider also the issue of system failures in providing the correct information to the agents. That is, a robust learning equilibrium is immune against both agents' failures and system failures. We apply this notion to repeated first-price auctions, and prove the existence

of robust learning equilibrium in that setting. Our proofs are constructive. This talk is based on the following papers:

"Learning in Resource Selection Games" - with Dov Monderer and Moshe Tennenholtz, AAAI-07

"Robust Learning Equilibrium" - with Dov Monderer and Moshe Tennenholtz, UAI-06

See also: Robust Learning Equilibrium - UAI 06

Item Pricing for Revenue Maximization in Combinatorial Auctions

Maria-Florina Balcan (CMU - Pittsburgh, USA)

Consider the problem of a retailer with various goods for sale, attempting to set prices to maximize revenue. If customers have separate valuations over the different goods, and these are known to the retailer, then the goods can be priced separately and the problem is not so difficult. However, when customers have valuations over sets of items, this becomes a combinatorial auction problem, and the problem becomes computationally hard even when valuations are fully known in advance. In this talk we present some simple randomized algorithms and mechanisms for a number of interesting cases of this problem, both in the limited and the unlimited supply settings.

This talk is based on joint work with Avrim Blum and Yishay Mansour.

Keywords: Item Pricing, Revenue Maximizing, Combinatorial Auctions

Joint work of: Blum, Avrim; Mansour, Yishai; Balcan, Maria-Florina

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1153>

Trading Networks with Price-Setting Agents

Larry Blume (Cornell University, USA)

In a wide range of markets, individual buyers and sellers often trade through intermediaries, who determine prices via strategic considerations.

Typically, not all buyers and sellers have access to the same intermediaries, and they trade at correspondingly different prices that reflect their relative amounts of power in the market.

We model this phenomenon using a game in which buyers, sellers, and traders engage in trade on a graph that represents the access each buyer and seller has to the traders. In this model, traders set prices strategically, and then buyers and sellers react to the prices they are offered. We show that the resulting game always has a subgame perfect Nash equilibrium, and that all equilibria lead to an efficient (i.e. socially optimal) allocation of goods. We extend these results

to a more general type of matching market, such as one finds in the matching of job applicants and employers. Finally, we consider how the profits obtained by the traders depend on the underlying graph. Roughly, a trader can command a positive profit if and only if it has an "essential" connection in the network structure, thus providing a graph-theoretic basis for quantifying the amount of competition among traders.

Our work differs from recent studies of how price is affected by network structure through our modeling of price-setting as a strategic activity carried out by a subset of agents in the system, rather than studying prices set via competitive equilibrium or by a truthful mechanism.

Joint work of: Blume, Larry; Easley, David; Kleinberg, Jon; Tardos, Eva

Full Paper:

<http://www.cs.cornell.edu/people/eva/traders.pdf>

Conversion Rates in Auctions for Sponsored Search

Liad Blumrosen (Microsoft - Mountain View, USA)

The generalized second-price auction (GSP) is used predominantly for sponsored search in search engines like Google, MSN-Live Search and Yahoo!. It has been shown by Edelman, Ostrovsky and Schwarz (2006) and by Varian (2006) that GSP admits an efficient pure Nash equilibrium. The model studied in Edelman et al and Varian assumed that all the clicks on the search engine ads gain the advertiser the same benefit. In practice, when an ad can be shown in one of several positions (a.k.a., slots) on the search results page, often the lower slots have higher acquisition rates. We study the implications of relaxing the identical acquisition rate assumption for GSP. In this case, we show that GSP does not always admit an efficient equilibrium anymore (neither pure nor mixed), even in the special case where ordering the advertisers by bid remains optimal. We show that when the bid space is discrete, an (inefficient) pure Nash equilibrium always exists. We characterize those equilibria and we also quantify their inefficiency.

Joint work of: Blumrosen, Liad; Hartline, Jason

Strategic Behavior in Multi-unit Assignment Problems: Theory and Evidence from Course Allocations

Estelle Cantillon (Université Libre de Brussels, B)

This paper analyses the assignment problem when agents have multi-unit demand. Applications include task assignment in a team, course allocation, sport drafts and any other allocation problem where money does not play a role in balancing supply and demand. There is no known allocation mechanism that is

ex-post efficient, strategyproof and minimally fair, and practical solutions must therefore trade off these different aspects. We study such a specific mechanism used at Harvard Business School to allocate courses to students. We argue that students in the HBS mechanism have an incentive to overreport their preferences for popular courses, that this incentive does not vanish with the size of the market and that it results in increased congestion. We confirm these predictions with detailed data on reported preferences and behavior in the HBS mechanism. We show that strategic behavior hurts students but that it might still be preferable to random serial dictatorship over course bundles, a strategyproof alternative.

Keywords: Course allocation, market design, assignment, multi-unit demand

Joint work of: Budish, Eric; Cantillon, Estelle

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1154>

Social Comparisons and Contributions to Online Communities: A Field Experiment on MovieLens

Yan Chen (University of Michigan, USA)

We explore the use of social comparison theory as a natural mechanism to increase contributions to an online movie recommendation community by investigating the effects of social information on user behavior in an online field experiment. We find that, after receiving behavioral information about the median user's total number of movie ratings, users below the median demonstrate a 530% increase in the number of monthly movie ratings, while those above the median decrease their monthly ratings by 62%. Movements from both ends converge towards the median, indicating conformity towards a newly-established social norm in a community where such a norm had been absent.

Furthermore, the social information has a more dramatic effect on those below the median, suggesting an interaction between conformity and competitive preferences. When given outcome information about the average user's net benefit score from the system, consistent with social preference theory, users with net benefit scores above average contribute 94% of the new updates in the database. In both treatments, we find a highly significant Red Queen Effect.

Keywords: Social comparison, conformity, public goods, embedded online field experiment

Joint work of: Chen, Yan; Harper, Maxwell; Konstan, Joseph; Li, Sherry

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1155>

Stochastically stable states in load balancing and congestion games

Christine Chung (University of Pittsburgh, USA)

Price of anarchy (POA) results assume players in a game will arrive at a Nash equilibrium. The Nash equilibrium solution concept has been criticized for unrealistically assuming that players are infinitely rational and have common knowledge of rationality. We show that POA upper bounds can apply in some games even when boundedly rational decision-making agents use simple selfish strategy selection rules to arrive at what economists Foster and Young termed "stochastically stable" states. We also show that in some games these stochastically stable states are the Nash equilibrium states that are socially optimal.

Keywords: Stochastic stability, congestion games, load balancing, evolutionary game theory, price of anarchy

Limited Verification of Identities to Induce False-Name-Proofness

Vincent Conitzer (Duke University, USA)

In open, anonymous environments such as the Internet, mechanism design is complicated by the fact that a single agent can participate in the mechanism under multiple identifiers. One way to address this is to design false-name-proof mechanisms, which choose the outcome in such a way that agents have no incentive to use more than one identifier. Unfortunately, there are inherent limitations on what can be achieved with false-name-proof mechanisms, and at least in some cases, these limitations are crippling. An alternative approach is to verify the identities of all agents. This imposes significant overhead and removes any benefits from anonymity. In this paper, we propose a middle ground. Based on the reported preferences, we check, for various subsets of the reports, whether the reports in the subset were all submitted by different agents. If they were not, then we discard some of them. We characterize when such a limited verification protocol induces false-name-proofness for a mechanism, that is, when the combination of the mechanism and the verification protocol gives the agents no incentive to use multiple identifiers. This characterization leads to various optimization problems for minimizing verification effort. We study how to solve these problems. Throughout, we use combinatorial auctions (using the Clarke mechanism) and majority voting as examples.

Keywords: Mechanism design, social choice, false-name-proofness, verifying identities, combinatorial auctions

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2007/1156>

See also: TARK 07 paper + new draft

Anonymity-Proof Voting Rules

Vincent Conitzer (Duke University, USA)

A (randomized, anonymous) voting rule maps any multiset of total orders of (aka. votes over) a fixed set of alternatives to a probability distribution over these alternatives. A voting rule f is neutral if it treats all alternatives symmetrically. It satisfies participation if no voter ever benefits from not casting her vote. It is false-name-proof if no voter ever benefits from casting additional (potentially different) votes. It is anonymity-proof if it satisfies participation and it is false-name-proof. We show that the class of anonymity-proof neutral voting rules consists exactly of the rules of the following form. With some probability $\alpha \in [0, 1]$, the rule chooses an alternative at random. With probability $1 - \alpha$, the rule first draws a pair of alternatives at random. If every vote prefers the same alternative between the two (and there is at least one vote), then the rule chooses that alternative. Otherwise, the rule flips a fair coin to decide between the two alternatives.

Keywords: Mechanism design, social choice, false-name-proofness, verifying identities, combinatorial auctions

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2007/1165>

See also: TARK 07 paper + new draft

Spectrum Auction Design

Peter Cramton (University of Maryland - College Park, USA)

Spectrum auctions are used by governments to assign and price licenses for wireless communication. The standard approach is the simultaneous ascending auction, in which many related lots are auctioned simultaneously in a sequence of rounds. I analyze the strengths and weaknesses of the approach. I then present a variation, the package clock auction, which addresses many of the problems of the simultaneous ascending auction while building on its strengths. The package clock auction is a simple dynamic auction in which bidders bid on packages of lots. Most importantly, the pricing rule and information policy are carefully tailored to mitigate gaming behavior. Truthful bidding is encouraged, which simplifies bidding and improves efficiency.

Keywords: Auctions, spectrum auctions, market design, package auction, clock auction, combinatorial auction

Computing Approximate Equilibria

Konstantinos Daskalakis (Univ. California - Berkeley, USA)

In view of the recent hardness results for mixed Nash equilibria, there has been increased interest in computing approximate equilibrium points, albeit progress has been slow. We will present recent developments on the subject, in particular algorithms which achieve constant factor approximations for 2-player games. Next, we will consider a very natural and important class of games, the anonymous games, in which, like in the stock market, a player is oblivious to the identities of the other players.

These games or variants are often appealing models for games of many players even so because the description size that they require is polynomial in the number of players -as opposed to exponential that a normal form game would require. We will present a polynomial time approximation scheme for the anonymous setting.

Auction Design with Avoidable Fixed Costs: An Experimental Approach

Wedad Elmaghraby (University of Maryland - College Park, USA)

Advances in information technology and computational power have opened the doors for auctioneers to explore a range of auction formats by considering varying degrees of bid expressivity and different payment rule, e.g., single price vs. discriminatory prices.

While it is clear that one can design more complicated auctions, it is still not clear if should do so and which auction parameters have the greatest impact on the performance on cost and efficiency.

The purpose of this paper is to gain some insight into this question, via analytical and experimental methods.

Keywords: Auctions, Experimental, Procurement, Synergies, Asymmetric Bidders

Joint work of: Elmaghraby, Wedad; Larson, Nathan

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2007/1157>

A Simple Combinatorial Auction

Jacob Goeree (CalTech - Pasadena, USA)

We introduce a new combinatorial auction format based on a simple, transparent pricing mechanism tailored for the hierarchical package structure proposed by Rothkopf, Pekeč, and Harstad (1998) to avoid computational complexity.

This combination provides the feedback necessary for bidders in multi-round auctions to discern winning bidding strategies for subsequent rounds and to coordinate responses to aggressive package bids. The resulting mechanism is tested in laboratory experiments involving a considerable number of licenses and many possible allocations. In particular, the experiments provide a “wind tunnel” test of three alternative auction formats that are being considered for a major upcoming FCC auction in the 700MHz range.

Keywords: FCC, Combinatorial Auctions, Experiments

Designing eBay’s new reputation system: Experimental and empirical evidence

Ben Greiner (Harvard Business School - Boston, USA)

EBay recently changed its feedback rules. We study the impact of this change on bidding behavior, seller trustworthiness and feedback behavior.

Joint work of: Bolton, Gary; Greiner, Ben; Ockenfels, Axel

Implementation and Money Burning

Jason D. Hartline (Microsoft - Mountain View, USA)

We consider the general problem of designing socially optimal, single round, sealed bid mechanisms when transfers made from the agents to the mechanism must be burnt (i.e., are a social loss). In these settings the socially optimal outcome is the one that maximizes the *marginal surplus*, that is, the sum of the agents’ valuations minus the agents’ payments (minus any social cost of the outcome). In the Bayesian setting where the agents’ valuations are drawn independently, but not necessarily identically, from a known distribution, we give a concise characterization of the mechanism that maximizes the expected marginal surplus. From this characterization we observe that the socially optimal way to allocate a single item to agents with i.i.d. valuations is, depending on the distribution, to either assign the item to an arbitrary agent or to run a second-price-like auction. Furthermore, for non-identical distributions that satisfy the monotone hazard rate condition, the socially optimal mechanism (for any given social cost function) is the one that chooses the allocation to maximize the expected social surplus ex ante and requires no payments.

Keywords: Optimal auction design, mechanism design

Farsightedly Stable Networks

Jean-Jacques Herings (Maastricht University, NL)

We propose a new concept, the pairwise farsightedly stable set, in order to predict which networks may be formed among farsighted players.

A set of networks G is pairwise farsightedly stable (i) if all possible pairwise deviations from any network $g \in G$ to a network outside G are deterred by the threat of ending worse off or equally well off, (ii) if there exists a farsightedly improving path from any network outside the set leading to some network in the set, and (iii) if there is no proper subset of G satisfying (i) and (ii). We show that a non-empty pairwise farsightedly stable set always exists and we provide a full characterization of unique pairwise farsightedly stable sets of networks. Contrary to other pairwise concepts, pairwise farsighted stability yields a Pareto dominating network, if it exists, as the unique outcome. Finally, we study the relationship between pairwise farsighted stability and other concepts such as the largest consistent set.

Keywords: Networks, farsighted, stability, pairwise, efficiency

Joint work of: Herings, Jean-Jacques; Mauleon, Ana; Vannetelbosch, Vincent

Full Paper:

<http://www.unimaas.nl/media/um-layout/fdewb/opmaak.htm?http://www.fdewb.unimaas.nl/algec/staff/framespages/Herings/herings.htm>

See also: Herings, P.J.J., A. Mauleon, and V. Vannetelbosch (2006), "Farsightedly Stable Networks"

On Revenue Equivalence in Truthful Mechanisms

Birgit Heydenreich (Maastricht University, NL)

The property of an allocation rule to be implementable in dominant strategies by a unique payment scheme is called revenue equivalence. In this paper we give a characterization of revenue equivalence based on a graph theoretic interpretation of the incentive compatibility constraints. The characterization holds for any (possibly infinite) outcome space and many of the known results about revenue equivalence are immediate consequences.

Keywords: Mechanism Design, Revenue Equivalence, Graph Theory

Joint work of: Heydenreich, Birgit; Müller, Rudolf; Uetz, Marc; Vohra, Rakesh

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1158>

Aggregation of binary evaluations

Ron Holzman (Technion - Haifa, IL)

We study a general aggregation problem in which a society has to determine its position (yes/no) on each of several issues, based on the positions of the members of the society on those issues.

There is a prescribed set of feasible evaluations, i.e., permissible combinations of positions on the issues. This framework for the theory of aggregation was introduced by Wilson and further developed by Rubinstein and Fishburn. Among other things, it admits the modelling of preference aggregation (where the issues are pairwise comparisons and feasibility reflects rationality), and of judgment aggregation (where the issues are propositions and feasibility reflects logical consistency). We characterize those sets of feasible evaluations for which the natural analogue of Arrow's impossibility theorem holds true in this framework.

Keywords: Aggregation, Arrow's impossibility, social choice, logic, judgment

Joint work of: Dokow, Elad; Holzman, Ron

Mechanism Design for Stochastic Optimization Problems

Samuel Ieong (Stanford University, USA)

We identify and address algorithmic and game-theoretic issues arising from welfare maximization in the well-studied two-stage stochastic optimization framework. In contrast, prior work in algorithmic mechanism design has focused almost exclusively on optimization problems without uncertainty. We show both positive results, by demonstrating a mechanism that implements the social welfare maximizer in *sequential ex post equilibrium*, and also negative results, by showing the impossibility of dominant-strategy implementation. We also introduce a novel generalization of the Fixed-Tree Multicast problem, for which welfare maximization is #P-hard. We develop a polynomial-time approximation scheme for the problem, and characterize a precise trade-off between the incentive guarantees and the degree of approximation.

An extended abstract of our work will appear in WINE 2007, and a journal version is under preparation.

Keywords: Stochastic Programming, Mechanism Design

Joint work of: Ieong, Samuel; So, Anthony Man-Cho; Sundararajan, Mukund

Secretary Problems and Online Auction Design

Nicole Immorlica (Microsoft Research - Redmond, USA)

In this talk, I will discuss the problem of online auction design, or auctions for bidders who arrive and depart over time. Maximizing welfare in such auctions is complicated by the fact that bids must be accepted or rejected at the moment they are submitted. It is known how the classic secretary problem introduced by Dynkin in 1963 can be used to design approximately welfare-maximizing auctions in a simple multi-unit auction setting. We show how the classic secretary problem

can be generalized to a combinatorial setting, and use this generalization to build mechanisms for a class of online combinatorial auctions. Parts of this talk are based on joint work with Moshe Babaioff, David Kempe, and Robert Kleinberg.

Joint work of: Babaioff, Moshe; Immorlica, Nicole; Kempe, David; Kleinberg, Robert

Learning spillover and analogy-based expectations: A multi-game experiment

Philippe Jehiel (University College London, GB)

We consider a multi-game interactive learning environment and ask ourselves whether long run behaviors in one game are affected by behaviors in the other, i.e. whether there are learning spillovers. Our main finding is that learning spillovers arise whenever the feedback provided to subjects about past play is not easily accessible game by game and thus subjects get a more immediate impression about aggregate distributions. In such a case, long run behaviors stabilize to an analogy-based expectation equilibrium (Jehiel 2005), thereby suggesting how one should broaden the notion of equilibrium to cope with learning spillovers.

Keywords: Experiment, analogy-based expectation

Joint work of: Jehiel, Philippe; Huck, Steffen; Rutter, Tom

Positive Externalities, Negative Externalities, and Optimal Scale

Ramesh Johari (Stanford University, USA)

We study a system where many identical customers use a single service.

Each customer experiences both positive externalities (a positive "network effect") and negative externalities (a "congestion effect") from other customers using the service. Such a model arises frequently in practice: application services on wireless networks are an example. We characterize the social optimum, where a social planner determines the usage level of each customer. We also characterize the Nash equilibrium achieved when the usage levels are determined by the customers themselves, in their self-interest. We study the ratio of the welfare in Nash equilibrium to that in the social optimum. We demonstrate that there is an "optimal scale", i.e., a number of customers at which this ratio is maximized; further, the optimal ratio is unity. We also show that this same optimal scale maximizes the Nash welfare of a single individual. We interpret our results in terms of club formation, and study the size of the club as the positive externality grows.

Joint work with Sunil Kumar, Stanford GSB.

An Experimental Investigation of Buyer Determined Procurement Auctions

Elena Katok (Penn State University, USA)

We present a laboratory study of buyer determined auctions in a setting in which non-monetary attributes are exogenous and influence buyer surplus. In buyer determined reverse auctions, the buyer does not commit to awarding a contract to the supplier who submits the lowest bid, but is instead free to select a winner based on the buyer surplus generated. We find that generally the sealed-bid auction format generates higher buyer surplus levels than the dynamic auction format. We also find that when the value of suppliers' non-monetary attributes is public information, dynamic auctions perform particularly poorly. A useful practical implication from our study is that buyers would prefer sealed bid auctions for procurement applications, especially when they have reasons to believe suppliers are aware of how much one another's non-monetary attributes are worth to the buyer.

Keywords: Bidding, Procurement, Reverse Auctions, Multi-Attribute Auctions, Behavioral Game Theory, Experimental Economics

Joint work of: Haruvy, Ernan; Katok, Elena

Online Pricing and Prophet Inequalities

Robert Kleinberg (Cornell University, USA)

In probability theory, a prophet inequality is a theorem which compares the expected payoff obtained by two parties each selecting k elements from the same sequence of random numbers: a prophet, who has foreknowledge of all the numbers in the sequence and may select any k of them, versus a gambler, who knows only the distribution from which the sequence is sampled and must make selections online. This talk will explore a relation between prophet inequalities and the analysis of optimal online pricing mechanisms for digital goods. We will present some new (and nearly optimal) multi-choice prophet inequalities, and we will apply them to understand the effect of different temporal incentive constraints on the expected revenue of the optimal dominant-strategy incentive-compatible online mechanism.

Joint work with MohammadTaghi Hajiaghayi, Tuomas Sandholm, and Ymir Vigfusson.

Approximability of Pricing Problems

Piotr Krysta (University of Liverpool, GB)

Suppose a seller wants to sell a finite collection of goods which can be available in limited or unlimited supply.

We have a set of potential customers and each customer specifies a single subset of goods she is interested in and the maximum price she is willing to pay for that subset. If the goods are the edges of a graph and customers are requesting to purchase paths in this graph, then we can think of the problem as pricing computer network connections or transportation links. We call such customers single-minded as they are interested in whole single subset of goods. The problem is to compute the prices for goods so as to maximize the overall revenue of the seller.

In another setting we will consider, called unit-demand, each customer also declares a single subset of goods together with non-zero budgets for each single good, and a ranking of all the goods the customer is interested in.

Once prices are fixed, each customer chooses to buy one of the goods she can afford based on some predefined selection rule, such as min-buying, max-buying, and rank-buying. Again, the objective is to find the prices of goods to maximize the revenue of the seller.

In this talk we will consider the approximability of such problems, and will discuss both approximation algorithms and non-approximability results for some variants of these problems.

Joint work of: Krysta, Piotr; Briest, Patrick

Reducing Costly Information Acquisition in Auctions

Kate Larson (University of Waterloo, CA)

Most auction research assumes that potential bidders have private information about their willingness to pay for an item. In reality, bidders often have to go through a costly information-gathering process in order to learn their valuations. Recent attempts at modelling this phenomena has brought to light complex strategic behaviour arising from information-gathering, and has shown that traditional approaches to auction and mechanism design are not able to overcome it.

Keywords: Auctions, Information Gathering

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1159>

Truthful Mechanism Design for Multi-Dimensional Scheduling via Cycle Monotonicity

Ron Lavi (Technion - Haifa, IL)

We consider the problem of makespan minimization on m unrelated machines in the context of algorithmic mechanism design, where the machines are the strategic players.

This is a multidimensional scheduling domain, and the only known positive result for makespan minimization in such a domain is an m -approximation truthful mechanism due to Nisan and Ronen [23]. (Even the existence of a nearoptimal truthful mechanism is not known.) We study a wellmotivated special case of this problem, where the processing time of a job on each machine may either be "low" or "high", and the low and high values are public and job-dependent.

This preserves the multidimensionality of the domain, and generalizes the restricted-machines (i.e., $p_j, 1$) setting in scheduling. We give a general technique to convert any c -approximation algorithm to a $3c$ -approximation truthful-in-expectation mechanism, and thus show that there exists a 3 -approximation truthful-in-expectation mechanism for this problem. This is one of the few known results that shows how to export approximation algorithms for a multidimensional problem into truthful mechanisms in a black-box fashion.

When the low and high values are the same for all jobs (which is still a multidimensional domain), we devise a deterministic 2 -approximation truthful mechanism. These are the first truthful mechanisms with non-trivial performance guarantees for a multidimensional scheduling domain.

Our constructions are novel in two respects. First, we do not utilize or rely on explicit price definitions to prove truthfulness; instead we design algorithms that satisfy cycle monotonicity. Cycle monotonicity [24] is a necessary and sufficient condition for truthfulness, which may be viewed as the parallel of value monotonicity for multidimensional domains. However, whereas value monotonicity has been used extensively and successfully to design truthful mechanisms in single-dimensional domains, ours is the first work that leverages cycle monotonicity in the multidimensional setting; our methods and analyses shed new light on its potential usefulness. Second, our randomized mechanisms are obtained by first constructing a fractional truthful mechanism for a fractional relaxation of the problem, and then converting it into a truthful-in-expectation mechanism. This builds upon a technique of [17], and shows the usefulness of fractional mechanisms in truthful mechanism design.

Joint work of: Lavi, Ron; Swamy, Chaitanya

Revenue Monotonicity in Combinatorial Auctions

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

Intuitively, one might expect that a seller's revenue from an auction weakly increases as the number of bidders grows, as this increases competition. However, it is known that for combinatorial auctions that use the VCG mechanism, a seller can sometimes increase revenue by dropping bidders. We investigate the extent to which this phenomenon can occur under other dominant-strategy combinatorial auction mechanisms. We show that such failures of "revenue monotonicity" are not limited to mechanisms that achieve efficient allocations. Instead, they can occur under any dominant-strategy direct mechanism that sets prices using

critical values, and that is maximal with respect to at least two bidders. We say that a mechanism is maximal with respect to bidder i if and only if the mechanism always chooses an allocation such that either (i) bidder i wins, or (ii) the allocation cannot be augmented to award bidder i a bundle that she values more than some predefined constant amount. Examples of such mechanisms include affine maximizers and mechanisms that always select weakly Pareto efficient allocations.

Keywords: Combinatorial Auctions, Revenue, Auction Theory, Game Theory

Joint work of: Rastegari, Baharak; Condon, Anne; Leyton-Brown, Kevin

Full Paper:

<http://cs.ubc.ca/~kevinlb/papers/2007-AAAI-RevenueMonotonicity.pdf>

Keeping the bad stuff out: structuring incentives for user behavior

Jeffrey MacKie-Mason (University of Michigan, USA)

Humans are “smart components” in a system, but cannot be directly programmed to perform; rather, their autonomy must be respected as a design constraint and incentives provided to induce desired behavior. Sometimes these incentives are properly aligned, and the humans don’t represent a vulnerability. But often, a misalignment of incentives causes a weakness in the system that can be exploited by clever attackers. Incentive-centered design tools help us understand these problems, and provide design principles to alleviate them. We describe incentive-centered design and some tools it provides. We provide a number of examples of security and related problems for which incentive centered design might be helpful. We elaborate with a general screening model that offers strong design principles for a class of security problems.

(The talk will be based on the attached paper, but I will also present another application, to user-contributed content.)

Keywords: Incentive-centered design, security, user-contributed content, social computing, game theory, economics

Joint work of: MacKie-Mason, Jeffrey; Wash, Rick

See also: Rick Wash and Jeffrey K. MacKie-Mason (2007),

Efficiency and Revenue Maximization in Dynamic Allocation Problems with Heterogenous Objects and Random Arrivals

Benny Moldovanu (Universität Bonn, D)

We study the dynamic allocation of several heterogenous objects to privately informed, randomly arriving agents.

First, we characterize all dynamic, incentive compatible allocation schemes. Then, we focus on two designer's goals:

1. efficiency and
2. revenue maximization,

in two timing scenarios:

1. the object must be sold by a deadline and
2. potentially infinite horizon with time discounting.

The solutions to these problems are given by complex systems of differential equations that can be seldom solved analytically. But, we use stochastic dominance arguments and insights from majorization theory in order to obtain qualitative results that apply to large classes of situations.

Mediated Equilibrium

Dov Monderer (Technion - Haifa, IL)

We define and analyze mediators for non-cooperative games. A mediator is a reliable entity, which can ask the agents for the right to play on their behalf, and is guaranteed to behave in a pre-specified way based on messages received from the agents.

However, a mediator can not enforce behavior; that is, agents can play in the game directly without the mediator's help. Hence, a mediator generates a new game for the players, the mediated game. We prove some general results about mediators, and mainly concentrate on the notion of strong mediated equilibrium, which is just a strong equilibrium at the mediated game. A strong equilibrium, which is a strategy profile immune to deviations by coalition rarely exists. In contrast, we show that desired behaviors, which are stable against deviations by coalitions, can be obtained using mediators in several classes of settings.

At the second part of the talk we define mediators for games with incomplete information, and construct such a mediator for position auctions.

The talk is based on two papers:

Dov Monderer, and Moshe Tennenholtz "Strong Mediated Equilibrium" AAAI-06
Itai Ashlagi, Dov Monderer, and Moshe Tennenholtz "Mediators for Position Auctions". EC-2007

Keywords: Mediators, Strong Equilibrium, Incomplete Information, Position Auctions

Joint work of: Ashlagi, Itai; Monderer, Dov; Tennenholtz, Moshe

Full Paper:

<http://ie.technion.ac.il/dov.phtml>

Strategy-proof assignment with a vanishing budget surplus

Hervé Moulin (Rice University - Houston, USA)

A VCG mechanism to assign p identical objects is feasible is cash transfers yield no deficit. The efficiency loss of such a mechanism is the worst ratio of budget surplus to efficient surplus. We compute the optimal efficiency loss for all n and p , when we also require Voluntary Participation as well as when we do not. Without the VP requirement, the optimal efficiency loss converges to zero uniformly in p , and exponentially fast if p is fixed. With the VP requirement asymptotic budget balance is only true if p is not larger than $n/2$.

Keywords: VCG mechanisms, assignment, asymptotic budget balance, worst case analysis

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1160>

Inefficiency of equilibria in query auctions with continuous valuations

Rudolf Müller (Maastricht University, NL)

We show that, when bidders have continuous valuations, any ex post equilibrium in an ex post individually rational query auction can only be ex post efficient when the running time of the auction is infinite for almost all realizations of valuations of the bidders. In contrast we show that, when we allow for inefficient allocations with arbitrarily small probability, there is a query auction (to be more specific, a bisection auction) that attains this level of approximate efficiency in equilibrium, while additionally the running time of the auction in equilibrium is finite for all realizations of valuations.

Keywords: Query auctions, ex post equilibrium, efficiency

Joint work of: Grigorieva, Elena; Herings, P. Jean-Jacques; Müller, Rudolf; Vermeulen, Dries

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2007/1161>

Efficient Online Mechanisms for Persistent, Periodically-Inaccessible Self-Interested Agents

David C. Parkes (Harvard University, USA)

We consider the problem of implementing a system-optimal decision policy in the context of self-interested agents with private state in an uncertain world.

Unique to our model is that we allow both persistent agents, with an agent having a local MDP model to describe how its local world evolves given actions by a center, and also periodically-inaccessible agents, with an agent unable to report information while inaccessible. We first review the dynamic-VCG mechanism of Bergemann and Valimaki (2006), which handles persistent agents. We offer an independent, simple proof of its correctness. We propose a generalized mechanism, dynamic-VCG#, to allow also for inaccessibility, and identify conditions for its correctness. In closing, we observe that the mechanism is equivalent to the earlier online-VCG mechanism of Parkes and Singh (2003) in a restricted model.

Keywords: Incentives, sequential decision making, (PO)MDPs, mechanism design

Joint work of: Cavallo, Ruggiero; Parkes, David C.; Singh, Satinder

Incentive Compatible Regression Learning

Ariel D. Procaccia (The Hebrew University of Jerusalem, IL)

We initiate the study of incentives in a general machine learning framework. We focus on a game theoretic regression learning setting where private information is elicited from multiple agents, which are interested in different distributions over the sample space. This conflict potentially gives rise to untruthfulness on the part of the agents. In the restricted but important case when distributions are degenerate, and under mild assumptions, we show that agents are motivated to tell the truth. In a more general setting, we study the power and limitations of mechanisms without payments. We finally establish that, in the general setting, the VCG mechanism goes a long way in guaranteeing truthfulness and efficiency.

Keywords: Machine learning, regression, mechanism design

Joint work of: Dekel, Ofer; Fischer, Felix; Procaccia, Ariel D.

Extended Abstract: <http://drops.dagstuhl.de/opus/volltexte/2007/1162>

Optimal Protocol Design for Networks with Selfish Users

Tim Roughgarden (Stanford University, USA)

We study how to design network protocols that minimize the worst-case efficiency loss caused by selfish end users. The goal is to identify the optimal protocol subject to natural implementation constraints. We illustrate this idea in the context of cost-sharing protocols for large networks. Our results build on a complete characterization of the budget-balanced protocols that always induce pure-strategy equilibria, which in turn is proved using a connection between potential games and the Shapley value.

Joint work of: Chen, Ho-Lin; Roughgarden, Tim; Valiant, Gregory

Manipulation-Resistant Recommender Systems

Rahul Sami (University of Michigan, USA)

An attacker can draw attention to items that don't deserve that attention by manipulating recommender systems. We describe an influence-limiting algorithm that can turn existing recommender systems into manipulation-resistant system. Honest reporting of ratings is the optimal strategy for raters who wish to maximize their influence. If an attacker can create only a bounded number of skills, the attacker can mislead only a small amount. However, the system eventually makes full use of information from honest, informative raters. In addition to the influence-limiting algorithm, a second contribution of the paper is a characterization of the information contribution of raters in terms of information-theoretic concepts of entropy reduction. This characterization is used both in setting influence limits and in analyzing the information loss created by those limits.

Joint work of: Sami, Rahul; Resnick, Paul

Signalling Preferences in Interviewing Markets

Michael A. Schwarz (Yahoo Research - Berkeley, USA)

The process of match formation in matching markets can be divided into three parts: information sharing, investments in information acquisition, and the formation of matches based on available information. The last stage where agents are assumed to know their preferences has been studied in seminal work of Gale and Shapley (1962), and a model of second stage costly information acquisition is introduced and studied in Lee and Schwarz (2007). This paper focuses on the first stage – information sharing – and examines mechanisms which allow workers to signal their preferences over matching partners prior to the assignment of interviews. The incentives of firms and workers vis-a-vis information revelation are partially aligned – all other things being equal, a worker prefers to have an interview with a firm that is high in his preference ranking and a firm prefers to invest in interviewing a worker who ranks a firm highly because such worker is more likely to accept a job if offered. However, the incentives are far from being perfectly aligned. For instance, if firms pay the full cost of interviewing, each worker would prefer to have as many interviews as possible, and in a world with bilateral communication no information is revealed as each workers would want to tell each firm that it is his first choice. But if communication is moderated through an intermediary or there is a restriction on the number of messages a worker can send, then cheap talk becomes informative. Currently existing market institutions that facilitate information exchange prior to interviewing are discussed.

Keywords: Cheap talk, job search, labor market, matching, interview assignment

Joint work of: Lee, Robin S.; Schwarz, Michael A.

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2007/1163>

Inapproximability of Congestion Games

Alexander Skopalik (RWTH Aachen, D)

A natural and convincing notion of approximation for equilibria in congestion games assumes that agents are ambivalent between strategies whose delay differs by less than a factor of α , for some $\alpha > 1$. In this talk, we show that computing an α -approximate equilibrium is PLS-complete, for any polynomial time computable $\alpha > 1$. Thus finding an approximate Nash equilibrium is as hard as finding an exact Nash equilibrium or solving any other problem in PLS. To our knowledge this is the first inapproximability result with respect to problems in PLS.

Keywords: Congestion games, local search, approximation

Trade-offs in Cost-sharing Mechanisms

Mukund Sundararajan (Stanford University, USA)

(Moulin'99) proposes a general framework to design truthful, budget-balanced mechanisms. These mechanisms are based on cross-monotonic cost-sharing methods. First, we discuss the efficiency loss of Moulin mechanisms in various combinatorial settings. Second, we show that cross-monotonicity is not necessary in the design of truthful, budget-balanced mechanisms.

Keywords: Cost-sharing mechanisms efficiency budget balance strategyproofness

Joint work of: Sundararajan, Mukund; Roughgarden, Tim; Mehta, Aranyak

Full Paper:

<http://portal.acm.org/citation.cfm?id=1250910.1250912&coll=GUIDE&dl=GUIDE,&type=series&idx=SERIES958&part=series&WantType=Proceedings&title=EC&CFID=15151515&CFTOKEN=618461#>

See also: @INPROCEEDINGstrade, AUTHOR = "T. Roughgarden and M. Sundararajan", TITLE = "New Trade-Offs in Cost-Sharing Mechanisms", BOOK-TITLE = "Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC)", YEAR = 2006, PAGES = "79–88"@inproceedings1250912, author = Aranyak Mehta and Tim Roughgarden and Mukund Sundararajan, title = Beyond moulin mechanisms, booktitle = EC '07: Proceedings of the 8th ACM conference on Electronic commerce, year = 2007, isbn = 978-1-59593-653-0, pages = 1–10, location = San Diego, California, USA, doi = <http://doi.acm.org/10.1145/1250910.1250912>, publisher = ACM Press, address = New York, NY, USA,

Measuring the Influence of Neighbors in Social Networks

Siddharth Suri (Cornell University, USA)

The spread of influence among individuals in a social network can be naturally modeled in a probabilistic framework, but it is challenging to reason about differences between various models as well as to relate these models to actual social network data. Here we consider two of the most fundamental definitions of influence, one based on a small set of snapshot observations of a social network and the other based on detailed temporal dynamics. The former is particularly useful because large-scale social network data sets are often available only in snapshots or crawls. The latter however provides a more detailed process model of how influence spreads. We study the relationship between these two ways of measuring influence, in particular establishing how to infer the more detailed temporal measure from the more readily observable snapshot measure. We validate our analysis using the history of social interactions on Wikipedia; the result is the first large-scale study to exhibit a direct relationship between snapshot and temporal models of social influence.

Keywords: Influence, Social Network, Process Model, Probabilistic Model

Joint work of: Cosley, Dan; Huttenlocher, Dan; Kleinberg, Jon; Lan, Xiangyang; Suri, Siddharth

Approximately Maximizing Efficiency and Revenue in Polyhedral Environments.

Eva Tardos (Cornell University, USA)

We consider a resource allocation game in convex environments.

Convex environments model a wide range of problems, including bandwidth sharing, some models of Adword auctions and general resource allocation. We extend the fair sharing mechanism for such resource allocation games, and show that our mechanism simultaneously creates approximately efficient allocations and approximately maximizes revenue.

Joint work of: Nguyen, Thanh; Tardos, Eva

Full Paper:

<http://www.cam.cornell.edu/~thanh/paper/game.pdf>

Programs and Mediators

Moshe Tennenholtz (Technion - Haifa, IL)

This talk is a brief introduction to our work on program equilibrium and routing mediators. We show that when agents' strategies are computer programs which are executed on a given computing device, one can exploit the Von-Neumann dual role of computer programs.

We show that this idea implies that the set of program equilibrium payoffs of a game coincides with the set of feasible and individually rational payoffs of it. Computers which run such computer programs can be seen as action mediators, who can act on behalf of agents based on their instructions. Action mediators significantly enrich the set of situations where we can obtain stability against correlated deviations by coalitions. Moreover, we introduce the study of routing mediators, in which the above mediators may possess information also about the behavior of agents who do not give the mediator the right of play. We study the relationships between the power of different routing mediators in establishing correlated strong equilibrium. We show a natural class of routing mediators that allow to implement fair and efficient outcomes as a correlated super-strong equilibrium in a very wide class of games.

This talk is based on the following two papers:

Tennenholtz, M., Program Equilibrium, Games and Economic Behavior, 2004.

Rozenfeld, O. and Tennenholtz, M., Routing Mediators, IJCAI 2007.

The interplay between incentive constraints and communication complexity in mechanism design

Timothy Van Zandt (INSEAD - Fontainebleau, F)

The curse of the Revelation Principle is that it leads us to focus on unrealistic mechanisms in which agents report all private information to the principal, who then makes all decisions centrally. This is not true when communication costs are introduced. Then partial communication of information, sequential back-and-forth conversations, and decentralization of decisions become part of optimal mechanisms. This paper looks at the interplay between the incentive constraints and the communication constraints. When can they be separated, so that e.g. one can separately study the incentive compatibility of a social choice rule and the minimum-cost communication protocol that realizes the social choice rule? When do incentive constraints increase the communication complexity of realizing a social choice rule?

Keywords: Mechanism design, communication complexity

License Auctions with Royalty Contracts for (Winners and) Losers

Elmar Wolfstetter (HU Berlin, D)

This paper revisits the licensing of a non-drastic process innovation by an outside innovator to a Cournot oligopoly. We propose a new mechanism that combines a restrictive license auction with royalty licensing. This mechanism is more profitable than standard license auctions, auctioning royalty contracts, fixed-fee

licensing, pure royalty licensing, and two-part tariffs. The key features are that royalty contracts are auctioned and that losers of the auction are granted the option to sign a royalty contract. Remarkably, combining royalties for winners and losers makes the integer constraint concerning the number of licenses irrelevant.

Keywords: Auctions, Patents, Licensing

Joint work of: Wolfstetter, Elmar; Giebe, Thomas

Full Paper: <http://drops.dagstuhl.de/opus/volltexte/2007/1151>

Complexity of Finding Correlated Equilibria for Extensive Games

Bernhard von Stengel (London School of Economics, GB)

This talk discusses open issues on the computation of correlated equilibria (CE) for extensive-form games.

One approach (joint work with Françoise Forges) is the "Extensive Form Correlated Equilibrium" (EFCE) where players get recommendations of moves that are revealed to them only when they have to make that move, rather than a full strategy (i.e., all moves) at the beginning of the game. Interestingly, the computational difficulty of EFCE seems related to the existence of pure-strategy equilibria.

Background:

In a game, a CE is a commonly known joint distribution on strategy profiles. A "mediator" draws a profile from the distribution and privately tells each player her strategy in the profile. The equilibrium property (defined by "incentive constraints") states that, given the posterior on the strategies recommended to the other players, it is optimal for the player to play the recommended strategy. If the distribution is a product distribution, this specializes to a Nash equilibrium.

CE are computationally simpler than Nash equilibria because the incentive constraints are linear in the joint probabilities for the strategy profiles. Existence follows from the existence of Nash equilibria, or by a two-step linear programming argument due to Nau/McCardle and Hart/Schmeidler. The latter has been used by Papadimitriou to find, in polynomial time and using the ellipsoid algorithm, a CE for COMPACTLY specified games, where the set of strategy profiles may be exponential (e.g., in a game of n players with binary choices), as long as product distributions on profiles can be payoff-evaluated in polynomial time (e.g., if the game is symmetric in the n players).

An EXTENSIVE GAME is a game tree with information sets that reflect imperfect information about the game state, a central model of game theory. Players have "perfect recall", so they always remember what they knew and did earlier. The game tree is a compact specification, but the number of strategies can be exponential because a strategy specifies a move for every information

set. This complicates finding CE because there could be a different posterior for each of the exponentially many possible recommendations of a pure strategy to a player. Finding a payoff-maximal CE, even for two players and games without chance moves, is NP-hard. However, finding just one CE may still be possible in polynomial time. An easy way out is to redefine the concept of CE. Namely, consider the AGENT FORM, where each information set is controlled by a separate "agent" of the respective player. An agent-form CE is found in polynomial time by Papadimitriou's approach for compact games. However, Nash equilibria of the agent form may no longer be Nash equilibria of the original game, due to the lack of coordination between the agents. We propose the EFCE where players are not decomposed into agents, but where, nonetheless, they only get recommendations of MOVES that are delayed until the player has to make the respective decision. The full set of moves (as in a strategy profile) is generated by the mediator at the beginning of the game. The set of EFCE can be described by polynomially many linear constraints for TWO-PLAYER games WITHOUT chance moves, using the known "sequence form" which is of the same size as the game tree. Interestingly, these are also games where existence of a pure-strategy Nash equilibrium can be decided in polynomial time, which is NP-hard if there are chance moves, even for zero-sum games.

Note on the current, preliminary set of slides (June 2007):

Many of those slides are devoted to explaining the above background concepts. The talk will focus on the new EFCE concept, and on open questions.

Keywords: Correlated equilibrium, computational complexity, game trees

Full Paper:

<http://www.maths.lse.ac.uk/Personal/stengel/bvs-publ.html#corr>