Algorithms as Mechanisms: The Price of Anarchy of Relax-and-Round

PAUL DÜTTING, London School of Economics
THOMAS KESSELHEIM, Max-Planck-Institut für Informatik
ÉVA TARDOS, Cornell University

Many algorithms, that are originally designed without explicitly considering incentive properties, are later combined with simple pricing rules and used as mechanisms. The resulting mechanisms are often natural and simple to understand. But how good are these algorithms as mechanisms? Truthful reporting of valuations is typically not a dominant strategy (certainly not with a pay-your-bid, first-price rule, but it is likely not a good strategy even with a critical value, or second-price style rule either). Our goal is to show that a wide class of approximation algorithms yields this way mechanisms with low Price of Anarchy.

The seminal result of Lucier and Borodin [2010] shows that combining a greedy algorithm that is an α-approximation algorithm with a pay-your-bid payment rule yields a mechanism whose Price of Anarchy is $O(\alpha)$. In this paper we significantly extend the class of algorithms for which such a result is available by showing that this close connection between approximation ratio on the one hand and Price of Anarchy on the other also holds for the design principle of *relaxation and rounding* provided that the relaxation is *smooth* and the rounding is *oblivious*.

We demonstrate the far-reaching consequences of our result by showing its implications for sparse packing integer programs, such as multi-unit auctions and generalized matching, for the maximum traveling salesman problem, for combinatorial auctions, and for single source unsplittable flow problems. In all these problems our approach leads to novel simple, near-optimal mechanisms whose Price of Anarchy either matches or beats the performance guarantees of known mechanisms.

Categories and Subject Descriptors: F.2 [**Theory of Computation**]: Analysis of Algorithms and Problem Complexity; J.4 [**Computer Applications**]: Social and Behavioral Sciences—*Economics*

Additional Key Words and Phrases: Algorithmic Game Theory, Price of Anarchy, Smoothness

1. INTRODUCTION

Mechanism design — or "reverse" game theory — is concerned with protocols, or mechanisms, through which potentially selfish agents interact with one another. The basic assumption is that the data is held by the agents, who may behave strategically. The goal is then to achieve outcomes that approximate the social optimum in a wide range of strategic equilibria.

The most sweeping positive result that one could possibly hope for in this context — with some professional bias of course — is a general reduction from mechanism design to algorithm design, showing that mechanism design is just as "easy" as algorithm

---

É. Tardos is supported in part by NSF grants CCF-0910940 and CCF-1215994, ONR grant N00014-08-1-0031, a Yahoo! Research Alliance Grant, and a Google Research Grant.
Author addresses: P. Dütting, Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, UK; T. Kesselheim, Max-Planck-Institut für Informatik, Campus E1 4, 66123 Saarbrücken, Germany; É. Tardos, Department of Computer Science, Cornell University, Gates Hall, Ithaca, NY 14853, USA.
Emails: p.d.duetting@lse.ac.uk, thomas.kesselheim@mpi-inf.mpg.de, and eva@cs.cornell.edu.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
EC'15, June 15–19, 2015, Portland, OR, USA. ACM 978-1-4503-3410-5/15/06 ...$15.00.
Copyright is held by the owner/author(s). Publication rights licensed to ACM.
http://dx.doi.org/10.1145/2764468.2764486

design. Specifically, one could hope that using algorithms as they are and charging bidders their respective bids yields mechanisms whose equilibria are close to optimal.

Why would this be appealing? Such a result would make the entire toolbox of algorithm design available to mechanism design, significantly broadening the tools currently available. It would also "hide" the incentives aspect from the designer, who would then no longer need to worry about possible manipulations through the agents. He could simply focus on the problem of computing near optimal solutions for the claimed input. Finally, the resulting mechanisms would enjoy a simplicity well beyond that found in most state-of-the art mechanisms.

Our goal in this paper is to identify general algorithm design principles that work well when used *as* mechanisms. We cannot expect this to be the case for all algorithms. Identifying algorithm design principles that automatically work well as mechanisms would, in some sense, give us the vocabulary to which we — as algorithm designers — should confine ourselves if we expect that our algorithm will be used in strategic environments. Specifically, it would equip us with the tools to design *simple* and *robust* mechanisms for such settings.

Lucier and Borodin [2010] showed that greedy algorithms have this property: Any equilibrium of a greedy algorithm that is an α-approximation algorithm is within $O(\alpha)$ of the optimal solution. Our main result is to show that the common design principle of *relaxation and rounding* as in — relax integer linear program to fractional domain, solve relaxation to optimality, and then convert into integer solution via randomized rounding — also preserves the approximation guarantee as Price of Anarchy guarantee provided that the relaxation is *smooth* and the rounding process is *oblivious* (more on this below).

This result has — as we show — far-reaching consequences in mechanism design: It leads to novel simple, yet near-optimal mechanisms for sparse packing integer programs, such as multi-unit auctions and generalized matching, for the maximum traveling salesman problem, for combinatorial auctions and for single source routing problems. In all cases we obtain Price of Anarchy bounds that match or beat known Price of Anarchy guarantees, or they are the first non-trivial guarantees for the respective problem.

1.1. Our Contributions

Our results concern the algorithmic blueprint of *relaxation and rounding* (see, e.g., [Vazirani 2001]). In this approach a problem Π is *relaxed* to a problem Π', with the purpose of rendering exact optimization computationally tractable. Having found the optimal relaxed solution x', another algorithm derives a solution x to the original problem. This process is typically called *rounding*. The best-known example are integer linear programs which are relaxed to fractional domains.

Many rounding schemes in text books as well as highly sophisticated ones are *oblivious*. That is, they do *not* require knowledge of the objective function. Up to this point, to the best of our knowledge, this property—though wide-spread—has never proven useful. In this paper, we show that oblivious rounding schemes preserve bounds on the Price of Anarchy. That is, applying an α-approximate oblivious rounding scheme on a problem with a Price of Anarchy bound β, the combined mechanism has Price of Anarchy at most $O(\alpha\beta)$.

We thus translate the relax-and-round approach from algorithm design into mechanism design: If we relax a problem into a problem with Price of Anarchy β and round the solution to the relaxed problem with an α-approximate oblivious rounding scheme, the resulting mechanism has a Price of Anarchy of $O(\alpha\beta)$. The bound does not only apply to Nash equilibria, but also extends to the Bayesian setting as well as to learning outcomes (coarse correlated equilibria). See Section 9 for discussion on the existence and computational complexity of finding such outcomes.

Main Result. Our main result leverages the power of the smoothness framework of Roughgarden [2009, 2012] and Syrgkanis and Tardos [2013].

At the heart of this framework is the notion of a (λ, μ)-smooth mechanism, where $\lambda, \mu \geq 0$. The main result is that a mechanism that is (λ, μ)-smooth achieves a Price of Anarchy of $\beta(\lambda, \mu) = \max(1, \mu)/\lambda$ with respect to a broad range of equilibrium concepts including learning outcomes. Furthermore, the simultaneous and sequential composition of (λ, μ)-smooth mechanisms is again (λ, μ)-smooth. Ideally, $\lambda = 1$ and $\mu \leq 1$ in which case this result tells us that all equilibria of the mechanism are socially optimal; otherwise, if $\lambda < 1$ or $\mu > 1$, then this result tells us which fraction of the optimal social welfare the mechanism is guaranteed to get at any equilibrium.

The other crucial ingredient to our main result is the notion of an *α-approximate oblivious rounding scheme*, where $\alpha \geq 1$. This is a (possibly randomized) rounding scheme for translating a solution x' to the relaxed problem Π' into a solution x to the original problem Π so that for all possible valuation profiles each agent is guaranteed to get, in expectation, a $1/\alpha$-fraction of the value that it would have had for the solution to the relaxed problem.

Clearly an α-approximate oblivious rounding scheme, when combined with optimally solving the relaxed problem, leads to an approximation ratio of α. We show that it also approximately preserves the Price of Anarchy of the relaxation. We focus on pay-your-bid mechanisms for concreteness. Our result actually applies to a broad range of mechanisms and can also be extended to include settings where the relaxation is not solved optimally; we discuss these extensions in Section 8.

THEOREM 1.1 (MAIN THEOREM, INFORMAL). *Consider problem Π and a relaxation Π'. Suppose the pay-your-bid mechanism M for Π is derived from the pay-your-bid mechanism M' for Π'. If M' is (λ, μ)-smooth, then M is $(\lambda/(2\alpha), \mu)$-smooth.*

COROLLARY 1.2. *The Price of Anarchy established via smoothness of mechanism M' of β translates into a smooth Price of Anarchy bound for mechanism M of $2\alpha\beta$ extending to both Bayesian Nash equilibria and learning outcomes.*

Our main theorem can be strengthened if the relaxation satisfies a slightly stronger smoothness condition, also parametrized by λ and μ, which all our application do. In this case we can show that the derived mechanism is $(\lambda/\alpha, \mu)$-smooth; and the corollary would read "a Price of Anarchy of β translates into a Price of Anarchy of $\alpha\beta$."

Applications. We demonstrate the far-reaching consequences of our result by applying it to a broad range of optimization problems. For each of these problems we show the existence of a smooth relaxation and the existence of an oblivious rounding scheme. We note that in all of our applications, it is important to use the relaxation to show smoothness of the problem. For example, optimally solving the original (integer) problem would give a very high Price of Anarchy.

Sparse Packing Integer Programs. The first problem we consider are multi-unit auctions with n bidders and m items, where bidders have unconstrained valuations. The underlying optimization problem has a natural LP relaxation, which we show is $(1/2, 1)$-smooth. Using the 8-approximate oblivious rounding scheme of [Bansal et al. 2010], our framework yields a constant PoA. This is quite remarkable as solving the integral optimization problem leads to a PoA that grows linearly in n and m.

We then consider the generalized assignment problem in which n bidders have unit-demand valuations for a certain amount of one of k services and allocations of services to bidders must respect the limited availability of each service. For this problem we also show $(1/2, 1)$-smoothness, and use the 8-approximate oblivious rounding scheme of [Bansal et al. 2010] to obtain a constant PoA.

Both these results are in fact special cases of a general result regarding sparse packing integer programs (PIP) that we show. Namely, the pay-your-bid mechanism that solves the canonical relaxation of a PIP with column sparsity d is $(1/2, d+1)$-smooth. Multi-unit auctions and the generalized assignment problem have $d = 1$; combinatorial auctions in which each bidder is interested in at most d items simultaneously have $d \geq 1$. For general PIPs the rounding scheme of [Bansal et al. 2010] is $O(d)$-approximate. We get a PoA of $O(d^2)$.

Maximum Traveling Salesman. Our second application is the maximization variant of the classic traveling salesman problem (max-TSP). We think of the problem as a game where each edge e has a value for being included, and the goal of the mechanism is to select a TSP of maximum total value. The classic algorithm for this problem is a 2-approximation [Fisher et al. 1979]. It proceeds by computing a cycle cover, dropping an edge from each cycle, and connecting the resulting paths in an arbitrary manner to obtain a solution. We prove this can be thought off as a 2-approximate oblivious rounding scheme and show, through a novel combinatorial argument, that the relaxation is $(1/2, 3)$-smooth. We thus obtain a Price of Anarchy of 12.

The best approximation guarantee for max-TSP is a $3/2$-approximation due to Kaplan et al. [2003]. The same approximation ratio is achieved by a (much simpler) algorithm of Paluch et al. [2012]. We show that this algorithm — just as the basic algorithm — can be interpreted as a relax-and-round algorithm. Generalizing the arguments for the basic algorithm to the (different) relaxation used in this interpretation, we show that this algorithm achieves a Price of Anarchy that is by a factor $3/4$ better than the Price of Anarchy of the basic algorithm.

These examples are especially interesting as they show how a seemingly combinatorial algorithm can be re-stated within our framework. They also represent the first non-trivial PoA bounds for this problem.

Combinatorial Auctions. We also consider the "canonical" mechanism design problem of combinatorial auctions in which valuations are restricted to come from a certain class. Our first result concerns fractionally subadditive, or XOS, valuations [Lehmann et al. 2006]. We show that the pay-your-bid mechanism for the canonical LP relaxation is $(1/2, 1)$-smooth. Using Feige's ingenious $e/(e-1)$-approximate oblivious rounding scheme [Feige 2009], our main result implies an upper bound on the Price of Anarchy of $4e/(e-1)$.

We then show how to extend this result to the recently proposed hierarchy of $\mathcal{MPH}$-k valuations [Feige et al. 2014]. Levels of the hierarchy correspond to the degree of complementarity in a given function. The lowest level $k = 1$ coincides with the class of XOS/fractionally subadditive valuations; the highest level $k = m$ can be shown to comprise all monotone valuation functions. We show that for $\mathcal{MPH}$-k valuations the LP relaxation is $(1/2, k+1)$ smooth. Together with the $O(k)$-approximate oblivious rounding scheme of [Feige et al. 2014] we obtain a Price of Anarchy of $\Theta(k^2)$.

These results nicely complement the recent work on "simple auctions" such as [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011; Feldman et al. 2013; Dütting et al. 2013; Roughgarden 2014], answering an open question of Babaioff et al. [2014] regarding the Price of Anarchy of *direct* mechanism based on approximation algorithms in these settings. The advantage of having a direct mechanism for this problem is that one can consider simple bidding strategies (such as bidding half the value) to establish the performance guarantees, whereas in indirect mechanisms such as combinatorial auctions with item bidding the computational effort is effectively shifted to the bidders.

Single Source Unsplittable Flow. The final problem that we consider are multi-commodity flow (MCF) problems with a single source (or target). In these problems we are given a capacitated, directed network and a set of requests consisting of a target (a source) and a demand, corresponding to requests of, say different information, held at the source. The goal is to maximize the total demand routed (or the total value of the demand routed), subject to feasibility. We assume each player has a demand for some flow to be routed from a shared source to a terminal specific to the player, and the player has a private value for routing this flow.

For this problem we show that the natural LP relaxation is $(1/2, 1)$-smooth. A $(1+\epsilon)$-approximate oblivious rounding scheme for high enough capacities is obtained through an adaptation of the "original" randomized rounding algorithm of [Raghavan 1988; Raghavan and Thompson 1987]. This yields a PoA of $2(1+\epsilon)$.

An interesting feature of this result is that the LP can be solved greedily through a variant of Ford-Fulkerson which allows us to exploit the known connection to smoothness [Lucier and Borodin 2010; Syrgkanis and Tardos 2013]. Crucially, the reference to these results has to be on the fractional level, as a greedy procedure on the integral level achieves a significantly worse approximation guarantee.

1.2. Related Work

Our work is closely related to the literature on so-called "back-box reductions", which has led to some of the most impressive results in algorithmic mechanism design (such as [Lavi and Swamy 2005; Briest et al. 2005; Dughmi and Roughgarden 2014; Dughmi et al. 2011; Babaioff et al. 2010, 2013]). This approach takes an algorithm, and aims to implement the algorithm's outcome via a game. To this end it typically modifies the algorithm and adds a sophisticated payment scheme. Our approach is different in that we consider an algorithm without any modification, introduce a simple payment rule, such as the "pay your bid" rule, and understand the expected outcomes of the resulting game.

Lavi and Swamy [2005] use *randomized meta rounding* [Carr and Vempala 2002] to turn LP-based approximation algorithms for packing domains into truthful-in-expectation mechanisms. Our result is similar in spirit as it demonstrates the implications of *obliviousness* for non-truthful mechanism design. The property that we need, however, is less stringent and shared by most rounding algorithms. Another important difference is that our approach is not limited to packing domains.

Briest et al. [2005] show how pseudo-polynomial approximation algorithms for single-parameter problems can be turned into a truthful fully polynomial-time approximation schemes (FPTAS). Dughmi and Roughgarden [2014] prove that every welfare-maximization problem that admits a FPTAS and can be encoded as a packing problem also admits a truthful-in-expectation randomized mechanism that is an FPTAS. Unlike our approach these approaches are limited to single-parameter problems, or to multi-parameter problems with packing structure.

Dughmi et al. [2011] present a general framework that also looks at the fractional relaxation of the problem. They show that if the rounding procedure has a certain property, which they refer to as *convex rounding*, then the resulting algorithm is truthful. They instantiate this framework to design a truthful-in-expectation mechanism for CAs with matroid-rank-sum valuations (which are strictly less general than submodular). The main difference to our work is that standard rounding procedures are often oblivious but typically not convex.

Babaioff et al. [2010, 2013] show how to transform a monotone or cycle-monotone algorithm into a truthful-in-expectation mechanism using a *single call* to the algorithm. The resulting mechanism coincides with the algorithm with high probability.

This work differs from ours in that it only applies to monotone or cycle-monotone algorithms.

By insisting on truthfulness, or truthfulness-in-expectation, as a solution concept, all these approaches face certain natural barriers to how good they can get (see, e.g., [Papadimitriou et al. 2008; Chawla et al. 2012]). In addition, they typically do not lead to simple, practical mechanisms. For example, despite running times technically being polynomial, these mechanisms require far more computational effort than standard approximation algorithms for the underlying optimization problem. In some cases, the reduction yields mechanisms in which the approximation guarantee is tight on every single instance (not only in the worst case). That is, even when the optimization problem is trivial, the mechanism sacrifices the solution quality for incentives.

1.3. Organization

We formally define our model in Section 2. Section 3 presents the meta-theorems with proofs. Sections 4 through 7 discuss applications. We only give proof sketches for these results here. Details can be found in the full version. Section 8 presents possible extensions of our framework. We conclude with a discussion or our results and its implications in Section 9.

2. PRELIMINARIES

Algorithm Design Basics. We consider maximization problems Π in which the goal is to determine a feasible outcome $x \in \Omega$ that maximizes total weight given by $w(x)$ for non-negative a weight function $w\colon \Omega \to \mathbb{R}_{\geq 0}$. A potentially randomized algorithm A receives the functions w as input and computes an output $A(w) \in \Omega$. The algorithm is an α-*approximation algorithm*, for $\alpha \geq 1$, if for all weights w, $\mathbf{E}\left[w(A(w))\right] \geq \frac{1}{\alpha} \cdot \max_{x \in \Omega} w(x)$.

We are interested in *relax-and-round algorithms*. These algorithms first relax the problem Π to Π' by extending the space of feasible outcomes to $\Omega' \supseteq \Omega$ and generalizing weight functions w to all $x \in \Omega'$. They compute an optimal solution $x' \in \Omega'$ to the relaxed problem. Then a solution $x \in \Omega$ of the original problem is derived based on $x' \in \Omega'$, typically via randomized rounding.

A rounding algorithm is *oblivious* if it does not require knowledge of the actual objective function w, beyond the fact that x' was optimized with respect to w. Formally, a rounding scheme is an α-*approximate oblivious rounding scheme* if, given some relaxed solution x', it computes a solution x such that for all w, $\mathbf{E}\left[w(x)\right] \geq \frac{1}{\alpha} w(x')$. Clearly, a relax-and-round algorithm based on an α-approximate oblivious rounding scheme is an α-approximation algorithm.

Mechanism Design Basics. Our results apply to general multi-parameter mechanism design problems Π in which agents $N = \{1, \ldots, n\}$ interact to select an element from a set Ω of outcomes. Each agent has a valuation function $v_i\colon \Omega \to \mathbb{R}_{\geq 0}$. We use v for the valuation profile that specifies a valuation for each agent, and v_{-i} to denote the valuations of the agents other than i. The quality of an outcome $x \in \Omega$ is measured in terms of its social welfare $\sum_{i \in N} v_i(x)$.

We consider direct mechanisms M that ask the agents to report their valuations. We refer to the reported valuations as bids and denote them by b. The mechanism uses outcome rule f to compute an outcome $f(b) \in \Omega$ and payment rule p to compute payments $p(b) \in \mathbb{R}_{\geq 0}$. Both the computation of the outcome and the payments can be randomized. We are specifically interested in *pay-your-bid mechanisms*, in which agents are asked to pay what they have bid on the outcome they get. In other words, in a pay-your-bid mechanism $M = (f, p)$, $p_i(b) = b_i(f_i(b))$. We assume that the agents have quasi-linear utilities and that they are risk neutral. That is, we assume that agent i's expected utility in mechanism $M = (f, p)$ is given by $u_i(b, v_i) = \mathbf{E}\left[v_i(f(b))\right] - \mathbf{E}\left[p_i(b)\right]$.

For the game-theoretic analysis we distinguish two settings. In the *complete information* setting agents know each others' valuations, and a potentially randomized bid profile b that may depend on v is a *mixed Nash equilibrium* if for all agents $i \in N$ and possible deviations b'_i that may depend on v, $\mathbf{E}_b[u_i(b, v_i)] \geq \mathbf{E}_{b'_i, b_{-i}}[u_i((b'_i, b_{-i}), v_i)]$. In the *incomplete information* setting valuations are drawn from independent distributions D_i, and each agent $i \in N$ knows its own valuation v_i and the distributions D_{-i} from which the other agents valuations are drawn. A *mixed Bayes-Nash equilibrium* is a potentially randomized bid profile b_i that may depend on this agent's valuation v_i and the distributions D_{-i} from which the other agents' valuations are drawn such that for all agents $i \in N$ and potential deviations b'_i which are also allowed to depend on v_i and D_{-i}, $\mathbf{E}_{b, v_{-i}}[u_i(b, v_i)] \geq \mathbf{E}_{b'_i, b_{-i}, v_{-i}}[u_i((b'_i, b_{-i}), v_i)]$.

Price of Anarchy. We evaluate the quality of mechanisms by their *Price of Anarchy*. The Price of Anarchy with respect to Nash equilibria (PoA) is the worst ratio between the optimal social welfare and the expected welfare in a mixed Nash equilibrium. Similarly, the Price of Anarchy with respect to Bayes-Nash equilibria (BPoA) is the worst ratio between the optimal expected social welfare and the expected welfare in a mixed Bayes-Nash equilibrium. Formally, define $\mathsf{NASH}(v)$ and $\mathsf{BNASH}(D)$ as the set of all mixed Nash and mixed Bayes Nash equilibria respectively. Then,

$$PoA = \max_{v} \max_{b \in \mathsf{NASH}(v)} \frac{\max_{x \in \Omega} \sum_{i \in N} v_i(x)}{\mathbf{E}\left[\sum_{i \in N} v_i(f(b))\right]} \text{ and } BPoA = \max_{D} \max_{b \in \mathsf{BNASH}(D)} \frac{\max_{x \in \Omega} \mathbf{E}\left[\sum_{i \in N} v_i(x)\right]}{\mathbf{E}\left[\sum_{i \in N} v_i(f(b))\right]}.$$

The Smoothness Framework. An important ingredient in our result is the following notion of a smooth mechanism of Syrgkanis and Tardos [2013]. A mechanism is (λ, μ)-*smooth* for $\lambda, \mu \geq 0$ if for all valuation profiles v and all bid profiles b there exists a possibly randomized strategy b'_i for every agent i that may depend on the valuation profile v of all agents and the bid b_i of that agent such that

$$\sum_{i \in N} \mathbf{E}\left[u_i((b'_i, b_{-i}), v_i)\right] \geq \lambda \cdot \max_{x \in \Omega} \sum_{i \in N} v_i(x) - \mu \cdot \sum_{i \in N} \mathbf{E}\left[p_i(b)\right] \quad .$$

THEOREM 2.1 (SYRGKANIS AND TARDOS [2013]). *If a mechanism is (λ, μ)-smooth and agents have the possibility to withdraw from the mechanism, then the expected social welfare at any mixed Nash or mixed Bayes-Nash equilibrium is at least $\lambda / \max(\mu, 1)$ of the optimal social welfare.*

As shown in [Syrgkanis and Tardos 2013], (λ, μ)-smoothness also implies a bound of $\max(\mu, 1)/\lambda$ on the Price of Anarchy for correlated equilibria, also known as learning outcomes. Furthermore, the simultaneous and sequential composition of multiple (λ, μ)-smooth mechanisms is again (λ, μ)-smooth. For details, on the precise definitions and statements beyond Nash equilibria, see [Syrgkanis and Tardos 2013].

In fact, our smoothness proofs show an even slightly stronger property, semi-smoothness as defined by [Caragiannis et al. 2015]: the deviation strategy b'_i only depends on the respective agent's valuation v_i, but not on the agent's bid b_i or the other agents' valuations v_{-i}. Therefore, the same Price of Anarchy bounds also apply to coarse correlated equilibria and Bayes-Nash equilibria with correlated types.

3. OBLIVIOUS ROUNDING AND SMOOTH RELAXATIONS

In this section, we show our main theorem. We consider mechanisms for a problem Π that are constructed as follows. First, one computes an optimal solution x' to a relaxed problem Π' that maximizes the declared welfare. That is, it maximizes $\sum_{i \in N} b_i(x')$. Afterwards, an α-approximate oblivious rounding scheme is applied to derive a feasible solution x to the original problem Π. Each bidder is charged $b_i(x)$, i.e., his declared value of this outcome.

THEOREM 3.1 (MAIN RESULT). *Consider problem Π and a relaxation Π'. Given a pay-your-bid mechanism $M' = (f', p')$ that is (λ, μ)-smooth where f' is an exact declared welfare maximizer for the relaxation Π'. Then a pay-your-bid mechanism $M = (f, p)$ for the original problem Π that is obtained from the relaxation through an α-approximate oblivious rounding scheme is $(\lambda/(2\alpha), \mu)$-smooth.*

In many applications, smoothness is shown by the deviation strategy of reporting half one's true value. First we show that, while generally the deviation strategy b'_i can be arbitrary, it is sufficient to consider only this deviation $b'_i = \frac{1}{2}v_i$. We exploit the fact that f' performs exact optimization.

LEMMA 3.2. *Given a pay-your-bid mechanism $M = (f, p)$ that is (λ, μ)-smooth where f is an exact declared welfare maximizer. Then M is $(\lambda/2, \mu)$-smooth for deviations to half the value. That is, for all bid vectors b and bids $b'_i = \frac{1}{2}v_i$ for all $i \in N$, $\sum_{i\in N} u_i((b'_i, b_{-i}), v_i) \geq \frac{\lambda}{2}OPT(v) - \mu \sum_{i\in N} p_i(b)$.*

PROOF. We first use (λ, μ)-smoothness of M. For any valuations, there have to be deviation bids fulfilling the respective conditions. So, in particular, let us pretend that each bidder i has valuation $\frac{1}{2}v_i$. By smoothness, there are bids b''_i against b such that

$$\sum_{i\in N} u_i\left((b''_i, b_{-i}), \frac{1}{2}v_i\right) \geq \lambda OPT\left(\frac{v}{2}\right) - \mu \sum_{i\in N} p_i(b) \quad . \tag{1}$$

The next step is to relate the sum of utilities $\sum_{i\in N} u_i((b'_i, b_{-i}), v_i) = \sum_{i\in N} \frac{1}{2}v_i(f_i(b'_i, b_{-i})) = \sum_{i\in N} b'_i(f_i(b'_i, b_{-i}))$ that agents with valuations v get in M when they unilaterally deviate from b to b'_i to the sum of utilities $\sum_{i\in N} u_i((b''_i, b_{-i}), \frac{1}{2}v_i)$ that they get in M with valuations $\frac{1}{2}v$ and unilateral deviations from b to b''_i.

The allocation function f optimizes exactly over its outcome space. Therefore, it can be used to implement a truthful mechanism $M^{\text{VCG}} = (f, p^{\text{VCG}})$ by applying VCG payments. As VCG payments are non-negative, we get

$$u_i((b'_i, b_{-i}), v_i) = \frac{1}{2}v_i(f(b'_i, b_{-i})) = b'_i(f(b'_i, b_{-i})) \geq b'_i(f(b'_i, b_{-i})) - p^{\text{VCG}}(b'_i, b_{-i}) \quad .$$

Observe that the latter term is exactly the utility bidder i receives in M^{VCG} if his valuation and bid is b'_i. As M^{VCG} is truthful, this term is maximized by reporting the true valuation. In other words, it can only decrease, if bidder i changes his bid to b''_i (keeping the valuation b'_i). That is,

$$u_i((b'_i, b_{-i}), v_i) \geq b'_i(f(b'_i, b_{-i})) - p^{\text{VCG}}(b'_i, b_{-i}) \geq b'_i(f(b''_i, b_{-i})) - p_i^{\text{VCG}}(b''_i, b_{-i}) \quad .$$

Finally, we use that p^{VCG} is no larger than p because VCG payments never exceed bids, i.e., $p_i^{\text{VCG}}(b''_i, b_{-i}) \leq b''_i(f(b''_i, b_{-i})) = p_i(b''_i, b_{-i})$. By furthermore changing b'_i back to $\frac{1}{2}v_i$, we get

$$u_i((b'_i, b_{-i}), v_i) \geq \frac{1}{2}v_i(f(b''_i, b_{-i})) - p_i(b''_i, b_{-i}) = u_i\left((b''_i, b_{-i}), \frac{1}{2}v_i\right) \quad .$$

Summing this inequality over all $i \in N$ and combining it with inequality (1), we get

$$\sum_{i\in N} u_i((b'_i, b_{-i}), v_i) \geq \lambda OPT\left(\frac{v}{2}\right) - \mu \sum_{i\in N} p_i(b) = \frac{\lambda}{2}OPT(v) - \mu \sum_{i\in N} p_i(b) \quad . \quad \square$$

It remains to show that smoothness of the relaxation for deviations to half the value, implies smoothness of the derived mechanism for the original problem. As it is often possible to directly show smoothness for deviations to half the value, we state the following stronger version of Theorem 3.1 for relaxations that are (λ, μ)-smooth for deviations to half the value.

Theorem 3.1 follows by first using Lemma 3.2 to argue that unconstrained (λ, μ)-smoothness of the relaxation implies $(\lambda/2, \mu)$-smoothness for deviations to half the value and then using Theorem 3.1′ to show that the derived mechanism is $(\lambda/(2\alpha), \mu)$-smooth.

THEOREM 3.1′ (STRONGER VERSION OF MAIN THEOREM). *If the pay-your-bid mechanism $M' = (f', p')$ that solves the relaxation Π' optimally is (λ, μ)-smooth for deviations to $b'_i = \frac{1}{2}v_i$, then the pay-your-bid mechanism $M = (f, p)$ for Π that is obtained from the relaxation through an α-approximate oblivious rounding scheme is $(\lambda/\alpha, \mu)$-smooth.*

PROOF. For any bid vector b, denote the utility of agent $i \in N$ under mechanism $M = (f, p)$ by $u_i(b, v) = v_i(f_i(b)) - p_i(b)$ and under mechanism $M' = (f', p')$ by $u'_i(b, v) = v_i(f'_i(b)) - p'_i(b)$.

For each bidder i, we consider the unilateral deviation by $b'_i = \frac{1}{2}v_i$. As M is a pay-your-bid mechanism, bidder i's utility when bidding b'_i against b_{-i} can be expressed by

$$\mathbf{E}\left[u_i((b'_i, b_{-i}), v_i)\right] = \mathbf{E}\left[v_i(f(b'_i, b_i)) - p_i(b'_i, b_{-i})\right] = \frac{1}{2}\mathbf{E}\left[v_i(f(b'_i, b_{-i}))\right] \quad .$$

Next we use that the outcome $f(b'_i, b_{-i})$ is derived from $f'(b'_i, b_{-i})$ by applying an α-approximate oblivious rounding scheme by considering the weight function in which $w_i = v_i$ for all i and concluding that $\mathbf{E}\left[v_i(f(b'_i, b_{-i}))\right] \geq \frac{1}{\alpha}v_i(f'(b'_i, b_{-i}))$. That is, for bidder i's utility, we get

$$\mathbf{E}\left[u_i((b'_i, b_{-i}), v_i)\right] \geq \frac{1}{2\alpha}v_i(f'(b'_i, b_{-i})) = \frac{1}{\alpha}u'_i((b'_i, b_{-i}), v_i) \quad ,$$

where the last step uses that M' is a pay-your-bid mechanism as well.

Next, we apply the fact that M' is (λ, μ)-smooth for deviations to $b'_i = \frac{1}{2}v_i$. We get for the sum of utilities in M that

$$\sum_{i \in N}\mathbf{E}\left[u_i((b'_i, b_{-i}), v_i)\right] \geq \frac{1}{\alpha}\sum_{i \in N}u'_i((b'_i, b_{-i}), v_i) \geq \frac{1}{\alpha}\left(\lambda OPT(v) - \mu\sum_{i \in N}p'_i(b)\right) \quad .$$

To bound the terms $p'_i(b)$, we use once more the fact that we are applying an α-approximate oblivious rounding scheme, this time to derive $f(b)$ from $f'(b)$ and considering the weight function in which $w_i = b_i$ for all i. This implies

$$p'_i(b) = b_i(f'_i(b)) \leq \alpha\mathbf{E}\left[b_i(f_i(b))\right] = \alpha\mathbf{E}\left[p_i(b)\right] \quad .$$

Overall, we get

$$\sum_{i \in N}\mathbf{E}\left[u_i((b'_i, b_{-i}), v_i)\right] \geq \frac{1}{\alpha}\lambda OPT(v) - \mu\sum_{i \in N}p_i(b) \quad ,$$

as claimed. □

4. SPARSE PACKING INTEGER PROGRAMS

In a sparse packing integer program (PIP) each bidder i can be served in K possible ways. The fact whether bidder i gets option k is represented by a binary variable

$x_{i,k} \in \{0,1\}$. Each bidder i can only get one option, that is $\sum_{k\in[K]} x_{i,k} \leq 1$ for each i. Furthermore, matrix A and vector c represent packing constraints between the bidders, requiring that $Ax \leq c$. Each bidder's valuation depends on the option that he is served by. That is v_i can be expressed as $v_i(x) = \sum_{k\in[K]} v_{i,k}x_{i,k}$. The goal is to find $\max \sum_{i\in N} v_i(x)$ subject to feasibility.

We consider the relaxation of this integer program in which the binary variables $x_{i,k} \in \{0,1\}$ are replaced with non-negative variables $x_{i,k} \geq 0$. The interpretation is that $x_{i,k}$ is a fractional allocation of option k to bidder i, and no bidder i can be assigned more than the fractional equivalent of one option. This relaxation is a LP and can therefore be solved in polynomial time.

The *column sparsity* d is the maximum number of non-zero entries in a single column of A. Formally, for each variable $x_{j,k}$, let $S_{j,k}$ be the set of constraints in A with a non-zero coefficient, that is, $S_{j,k} = \{\ell \mid A_{\ell,j,k} \neq 0\}$. Now $d = \max_{j,k}|S_{j,k}|$. Examples with $d = 1$ are multi unit-auctions with unconstrained valuations or unit demand auctions, where each player wants at most one item, possibly with player dependent capacity constraints, like makespan constraints in a generalized assignment problem; or more generally, combinatorial auctions in which each bidder is interested in bundles of at most d items are an example with $d \geq 1$.

THEOREM 4.1. *There is an oblivious rounding based, pay-your-bid mechanism for d-sparse packing integer programs that achieves a Price of Anarchy of* 32 *for* $d = 1$ *and of* $16d(d+1)$ *for general* d.

PROOF SKETCH. An $8d$-approximate oblivious rounding scheme is available from [Bansal et al. 2010]. We show that the LP relaxation of a d-sparse PIP is $(1/2, d+1)$-smooth for deviations to $b'_i = 1/2v_i$. Theorem 3.1′ then implies the Price of Anarchy guarantee.

To establish smoothness we first show that the mechanism is $(\frac{1}{2}, \mu)$-smooth for deviations to $b'_i = \frac{1}{2}v_i$ with μ defined as follows. Denoting the optimal declared welfare for capacity vector c and bid vector b by $W^b(c)$, we define $\mu > 0$ to be the smallest value such that for all feasible allocations $\bar{x}$, $\sum_{i\in N}(W^{b_{-i}}(c) - W^{b_{-i}}(c - A(\bar{x}_i, 0))) \leq \mu \sum_{i\in N} W^b(c)$.

We then show that $\mu = (d+1)$ is a valid solution to this problem through the following scaling argument: Construct from $W^b(c)$ a feasible allocation $\hat{x}^{-i}$ for capacities $c - A(\bar{x}_i, 0)$ by setting $\hat{x}^{-i}_{j,k} = (1 - \delta^i_{j,k})\hat{x}_{j,k}$, where $\delta^i_{j,k} = \max_{\ell\in S_{j,k}} \frac{(A(\bar{x}_i,0))_\ell}{c_\ell}$. Then for all j and k, $\sum_i \delta^i_{j,k} \leq \sum_i \sum_{\ell\in S_{j,k}} \frac{(A(\bar{x}_i,0))_\ell}{c_\ell} = \sum_{\ell\in S_{j,k}} \sum_i \frac{(A(\bar{x}_i,0))_\ell}{c_\ell} \leq |S_{j,k}| \leq d$ and therefore $\sum_{i\neq j, i\in N}(1 - \delta^i_{j,k}) \geq n - d - 1$. It follows that $\sum_{i\in N} W^{b_{-i}}(c - A(\bar{x}_i, 0)) \geq \sum_{j\in N} \sum_{i\neq j, i\in N} \sum_k b_{j,k} \left(1 - \delta^i_{j,k}\right) \hat{x}_{j,k} \geq (n-d-1)W^b(c)$, and therefore $\sum_{i\in N}(W^{b_{-i}}(c) - W^{b_{-i}}(c - A(\bar{x}_i, 0))) \leq (d+1)W^b(c)$. □

In stark contrast, as we show in the full version, the mechanism that solves the integral problem optimally has an unbounded Price of Anarchy even when $d = 1$.

5. SINGLE SOURCE UNSPLITTABLE FLOW

We consider the single source weighted unsplittable multi-commodity flow problem in which we are given a graph $G = (V, E)$ with edge capacities c_e for each edge $e \in E$. All bidders share a source node s and each bidder i has a sink node t_i. He asks for a path connecting s and t_i fulfilling his demand d_i. His value for this is v_i, and he has no value for less flow than his demand. We assume that the sink t_i and demand d_i for each player is common knowledge, so the player's bid is a claimed value, which will be denoted by b_i.

Let $\mathcal{P}_i$ be the paths connecting s and t_i. For each $P \in \mathcal{P}_i$, we have a variable $f_{i,P}$ denoting the amount of flow along path P. The problem requires single path routing, that is, all the d_i flow satisfying player i's demand must be carried by a single path. We use the following standard LP relaxation that maximizes $\sum_{i\in N}\sum_{P\in\mathcal{P}_i} b_i f_{i,P}$ subject to $\sum_{i\in N}\sum_{P\in\mathcal{P}_i: e\in P} f_{i,P} \le c_e$ for all $e \in E$ and $\sum_{P\in\mathcal{P}_i} f_{i,P} \le d_i$ for all $i \in N$.

Substituting $f_{i,P}$ by $d_i\bar{x}_{i,P}$, we get an LP formulation in the spirit of Section 4. However, this LP is not necessarily sparse, as the column sparsity d corresponds to the maximum path length. Nevertheless we are able to establish the following theorem.

THEOREM 5.1. *Suppose the minimum edge capacity is by a logarithmic factor larger than the maximum demand, i.e.,* $\min_{e\in E} c_e \ge c\epsilon^{-1}\log|E|\max_{i\in N} d_i$ *for some* $\epsilon > 0$ *and an appropriate constant* $c > 0$. *Then there is an oblivious rounding based, pay-your-bid mechanism for the single source unsplittable flow problem with Price of Anarchy at most* $2(1+\epsilon)$.

PROOF SKETCH. For the setting considered here Raghavan and Thompson [Raghavan 1988; Raghavan and Thompson 1987] present a $(1+\epsilon)$-approximate oblivious rounding scheme. For smoothness we show that the LP relaxation can be solved exactly using a loser-independent greedy heuristic: route flow using augmenting path giving priority to terminals with higher b_i/d_i value. This lets us leverage the known connection between loser-independence and smoothness [Lucier and Borodin 2010; Syrgkanis and Tardos 2013]. Specifically, we show $(1/2, 1)$-smoothness for deviations to $b'_i = \frac{1}{2}v_i$. The Price of Anarchy bound then follows by Theorem 3.1′. □

Importantly, the reference to greedy is on the fractional level, as the greedy algorithm for the integral problem can be as bad as an $\Omega(\sqrt{|E|})$ (see [Kleinberg 1996]). Also, as we show in the full version solving the integral problem optimally again leads to an unbounded PoA, even if there is a single source, a single target and just one unit capacity edge between the two.

6. MAX-TSP

In the asymmetric maximization version of the traveling salesperson problem, one is given a complete digraph $G = (V, E)$ with non-negative weights $(w_e)_{e\in E}$. Players are the edges with value w_e for being selected, and the mechanism aims to select a Hamiltonian cycle C that maximizes $\sum_{e\in C} w_e$. We show how existing combinatorial algorithms for this problem can be interpreted as relax-and-round algorithms, and derive the following theorem.

THEOREM 6.1. *There is a pay-your-bid mechanism for the maximum traveling salesman problem based on oblivious rounding that achieves a Price of Anarchy of* 9.

PROOF SKETCH. The algorithm of Fisher et al. [1979] "rounds" the problem of finding a Hamiltonian cycle to the problem of finding a cycle cover as follows: It first determines a maximum-weight cycle cover, then drops the minimum weight edge from each cycle, and connects the resulting vertex-disjoint paths in an arbitrary way to obtain a Hamiltonian cycle. This can be turned into a oblivious rounding scheme by dropping an random edge from each cycle. The resulting rounding scheme is a 2-approximation as each edge in the cycle cover is included in the final outcome with probability at least one half.

To be able to apply Theorem 3.1′ and obtain the Price of Anarchy bound it remains to show that the pay-your-bid mechanism that finds a cycle cover is $(1/2, 3)$-smooth. Our proof of this follows a similar pattern as the smoothness proof in Theorem 4.1, but unlike this proof it is not a scaling argument. Rather for edges e not in the optimal cycle cover, we consider the social cost of adding e to the solution. The key idea is to show that for any cycle cover C' this total net social cost is bounded by 3 times the optimal declared welfare by modifying the optimal cycle cover to force each of the edges $e \in C'$ into the solution one-by-one.

In the full version, we show how to strengthen this result to a Price of Anarchy of 9 by using the algorithm of Paluch et al. [2012] instead. □

7. COMBINATORIAL AUCTIONS

In this section, we consider combinatorial auctions (CAs). In a CA, m items are sold to n bidders. Each item is allocated to at most one bidder and each bidder i has a valuation $v_i(S)$ for the subset $S \subseteq [m]$ of items he receives. The canonical relaxation as a *configuration LP* uses variables $x_{i,S} \in [0,1]$ representing the fraction that bidder i receives of set S. The goal is to maximize $\sum_{i \in N} \sum_{S \subseteq [m]} b_i(S) \bar{x}_{i,S}$ s.t. $\sum_{i \in N} \sum_{S : j \in S} x_{i,S} \leq 1$ for all $j \in [m]$ and $\sum_S x_{i,S} \leq 1$ for all $i \in N$.

For arbitrary valuation functions, only very poor approximation factors can be achieved for the optimization problem. Therefore, we focus on XOS or fractionally subadditive valuations. That is, each valuation function v_i has a representation of the following form. There are values $v_{i,j}^\ell \geq 0$ such that $v_i(S) = \max_\ell \sum_{j \in S} v_{i,j}^\ell$. Feige et al. [2014] very recently generalized the class of XOS functions to $\mathcal{MPH}$-k, where XOS is precisely the case $k = 1$. A valuation function v_i belongs to class $\mathcal{MPH}$-k if there are values $v_{i,T}^\ell \geq 0$ such that $v_i(S) = \max_\ell \sum_{T \subseteq S, |T| \leq k} v_{i,T}^\ell$.

THEOREM 7.1. *There is a pay-your-bid mechanism that is based on oblivious rounding and achieves a Price of Anarchy of $4\frac{e}{e-1}$ for XOS-valuations and of $O(k^2)$ for $\mathcal{MPH}$-k-valuations.*

PROOF SKETCH. For general $\mathcal{MPH}$-k-valuations Feige et al. [2014] present a $O(k+1)$-approximate rounding scheme; a better constant of $\frac{e}{e-1}$ for the special case XOS can be achieved via the rounding scheme described in Feige [2009]. Both schemes are oblivious. Regarding smoothness we show that the configuration LP is $(1/2, k+1)$-smooth for deviations to $b_i' = \frac{1}{2} v_i$. The proof is analogous to the proof of Theorem 4.1. It is based on considering the net welfare loss of switching a player to his/her allocation in a different solution. We show that for any alternate allocation, the sum of the net welfare losses over all players can be bounded by $(k+1) W^b(A)$, which is $k+1$ times the maximum total declared value of any solution. The claimed Price of Anarchy bounds then follow from Theorem 3.1′. □

8. EXTENSIONS

Throughout this paper, we focused on pay-your-bid payment schemes. However, all of our results generalize to payment schemes that use arbitrary non-negative payments which are upper bounded by the respective bid. In this case, we resort to weak smoothness [Syrgkanis and Tardos 2013]. In our statements (λ, μ)-smoothness would be replaced by weak $(\lambda, 0, \mu)$-smoothness. Considering equilibria without overbidding, i.e., always $b_i(x) \leq v_i(x)$, this implies a PoA bound of $(1+\mu)/\lambda$.

Furthermore, Theorem 3.1 also holds if f' is not an exact declared welfare maximizer, but only allows implementation as a truthful mechanism. The interesting consequence is that it might make sense to only approximately solve the relaxation if this improves the smoothness guarantees. For example, a packing LP can be solved using the fractional-overselling mechanism in [Hoefer et al. 2013], which was originally introduced in [Krysta and Vöcking 2012]. The allocation rule is an $O(\log n + \log L)$-approximation for any packing LP with n bidders and L constraints between bidders. It allows implementation as a truthful mechanism but it is also a greedy algorithm in the sense of [Lucier and Borodin 2010]. Therefore, the respective pay-your-bid mechanism is $(\frac{1}{O(\log n + \log L)}, 1)$-smooth. This means that combining this algorithm with any α-approximate oblivious rounding scheme for the respective packing LP, we get a pay-your-bid mechanism with Price of Anarchy at most $O(\alpha(\log n + \log L))$.

Finally, Carr and Vempala [2002] introduced randomized metarounding, which is a technique to derive oblivious rounding schemes from non-oblivious ones. Lavi and Swamy [2005] used this result to construct truthful mechanisms. However, they additionally need a packing structure. As in our case oblivious rounding is enough, any rounding scheme derived from the original version in [Carr and Vempala 2002] is enough for our considerations.

9. DISCUSSION

In this paper we have shown that algorithms that follow the relax-and-round paradigm and whose rounding is oblivious have a very desirable property: Namely, if the rounding scheme is α-approximate and the relaxation has a Price of Anarchy via smoothness of $O(\beta)$, then the resulting relax-and-round mechanism has a Price of Anarchy of $O(\alpha\beta)$ provable via smoothness.

Two aspects that we did not touch upon are equilibrium existence and the computational complexity of computing an equilibrium. The former is particularly relevant for pure equilibrium concepts, such as pure Nash equilibria or pure Bayes-Nash equilibria. The latter has been shown to be a problem for Bayes-Nash equilibria in simultaneous first-price auctions [Cai and Papadimitriou 2014].

Our foremost intended application is to repeated settings, where regret minimization converges to a coarse correlated equilibrium in polynomial-time. In fact, we do not even need vanishingly small regret — we only need that agents have no regret for deviations to half their value. This argument readily applies to settings of incomplete information, showing near-optimal system performance even out of equilibrium. We consider the availability of such simple fall-back strategies as a major advantage of direct mechanisms over indirect mechanisms, where bidders typically have to solve a non-trivial problem to figure out good strategies.

REFERENCES

BABAIOFF, M., KLEINBERG, R., AND SLIVKINS, A. 2013. Multi-parameter mechanisms with implicit payment computation. In *Proc. of the 13th Conference on Economics and Computation*. 35–52.

BABAIOFF, M., KLEINBERG, R. D., AND SLIVKINS, A. 2010. Truthful mechanisms with implicit payment computation. In *Proc. of the 11th Conference on Electronic Commerce*. 43–52.

BABAIOFF, M., LUCIER, B., NISAN, N., AND LEME, R. P. 2014. On the efficiency of the walrasian mechanism. In *Proc. of the 15th Conference on Economics and Computation*. 783–800.

BANSAL, N., KORULA, N., NAGARAJAN, V., AND SRINIVASAN, A. 2010. On k-column sparse packing programs. In *Proc. of the 14th Conference on Integer Programming and Combinatorial Optimization*. 369–382.

BHAWALKAR, K. AND ROUGHGARDEN, T. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *Proc. of the 22nd Symposium on Discrete Algorithms*. 700–709.

BRIEST, P., KRYSTA, P., AND VÖCKING, B. 2005. Approximation techniques for utilitarian mechanism design. In *Proc. of the 37th Symposium on Theory of Computing*. 39–48.

CAI, Y. AND PAPADIMITRIOU, C. H. 2014. Simultaneous bayesian auctions and computational complexity. In *Proc. of the 15th Conference on Electronic Commerce*. 895–

910.
CARAGIANNIS, I., KAKLAMANIS, C., KANELLOPOULOS, P., KYROPOULOU, M., LUCIER, B., LEME, R. P., AND TARDOS, É. 2015. On the efficiency of equilibria in generalized second price auctions. *Journal of Economic Theory*, 343–388.
CARR, R. D. AND VEMPALA, S. 2002. Randomized metarounding. *Random Structures and Algorithms 20,* 3, 343–352.
CHAWLA, S., IMMORLICA, N., AND LUCIER, B. 2012. On the limits of black-box reductions in mechanism design. In *Proc. of the 44th Symposium on Theory of Computing*. 435–448.
CHRISTODOULOU, G., KOVÁCS, A., AND SCHAPIRA, M. 2008. Bayesian combinatorial auctions. In *Proc. of the 35th International Colloquium on Automata, Languages and Programming*. 820–832.
DUGHMI, S. AND ROUGHGARDEN, T. 2014. Black-box randomized reductions in algorithmic mechanism design. *SIAM Journal on Computing 43,* 1, 312–336.
DUGHMI, S., ROUGHGARDEN, T., AND YAN, Q. 2011. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *Proc. of the 43rd Symposium on Theory of Computing*. 149–158.
DÜTTING, P., HENZINGER, M., AND STARNBERGER, M. 2013. Valuation compressions in VCG-based combinatorial auctions. In *Proc. of the 9th Conference on Web, Internet and Network Economics*. 146–159.
FEIGE, U. 2009. On maximizing welfare when utility functions are subadditive. *SIAM Jornal on Computing 39,* 1, 122–142.
FEIGE, U., FELDMAN, M., IMMORLICA, N., IZSAK, R., LUCIER, B., AND SYRGKANIS, V. 2014. A unifying hierarchy of valuations with complements and substitutes. *Electronic Colloquium on Computational Complexity 21*, 103.
FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2013. Simultaneous auctions are (almost) efficient. In *Proc. of the 45th Symposium on Theory of Computing*. 201–210.
FISHER, M., NEMHAUSER, G., AND WOLSEY, L. 1979. An analysis of approximations for finding a maximum weight hamiltonian circuit. *Operations Research 27,* 4, 799–809.
HOEFER, M., KESSELHEIM, T., AND VÖCKING, B. 2013. Truthfulness and stochastic dominance with monetary transfers. In *Proc. of the 14th Conference on Economics and Computation*. 567–582.
KAPLAN, H., LEWENSTEIN, M., SHAFRIR, N., AND SVIRIDENKO, M. 2003. Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs. In *Proc. of the 44th Symposium on Foundations of Computer Science*. 56–65.
KLEINBERG, J. 1996. *Approximation Algorithms for Disjoint Paths Problems*. Ph.D. Thesis, Department of EECS, MIT.
KRYSTA, P. AND VÖCKING, B. 2012. Online mechanism design (randomized rounding on the fly). In *Proc. of the 39th International Colloquium on Automata, Languages and Programming*. 636–647.
LAVI, R. AND SWAMY, C. 2005. Truthful and near-optimal mechanism design via linear programming. In *Proc. of the 46th Symposium on Foundations of Computer Science*. 595–604.
LEHMANN, B., LEHMANN, D. J., AND NISAN, N. 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior 55,* 2, 270–296.
LUCIER, B. AND BORODIN, A. 2010. Price of anarchy for greedy auctions. In *Proc. of the 21st Symposium on Discrete Algorithms*. 537–553.
PALUCH, K. E., ELBASSIONI, K. M., AND VAN ZUYLEN, A. 2012. Simpler approximation of the maximum asymmetric traveling salesman problem. In *Proc. of the 29th International Symposium on Theoretical Aspects of Computer Science*. 501–506.

PAPADIMITRIOU, C. H., SCHAPIRA, M., AND SINGER, Y. 2008. On the hardness of being truthful. In *Proc. of the 49th Symposium on Foundations of Computer Science*.

RAGHAVAN, P. 1988. Probabilistic construction of deterministic algorithms: Approximating packing integer programs. *Journal on Computer Systems Sciences 37,* 2, 130–143.

RAGHAVAN, P. AND THOMPSON, C. 1987. Randomized rounding: a technique for provably good algorithms and algorithmic proofs. *Combinatorica 7*, 365–374.

ROUGHGARDEN, T. 2009. Intrinsic robustness of the price of anarchy. In *Proc. of the 41st Symposium on Theory of Computing*. 513–522.

ROUGHGARDEN, T. 2012. The price of anarchy in games of incomplete information. In *Proc. of the 13th Conference on Electronic Commerce*. 862–879.

ROUGHGARDEN, T. 2014. Barriers to near-optimal equilibria. In *Proc. of the 55th Symposium on Foundations of Computer Science*. 71–80.

SYRGKANIS, V. AND TARDOS, É. 2013. Composable and efficient mechanisms. In *Proc. of the 45th Symposium on Theory of Computing*. 211–220.

VAZIRANI, V. V. 2001. *Approximation Algorithms*. Springer-Verlag New York, Inc., New York, NY, USA.