Algorithms against Anarchy: Understanding Non-Truthful Mechanisms

PAUL DÜTTING, London School of Economics
THOMAS KESSELHEIM, Max-Planck-Institut für Informatik

The algorithmic requirements for dominant strategy incentive compatibility, or truthfulness, are well understood. Is there a similar characterization of algorithms that when combined with a suitable payment rule yield near-optimal welfare in all equilibria?

We address this question by providing a tight characterization of a (possibly randomized) mechanism's Price of Anarchy provable via smoothness, for single-parameter settings. The characterization assigns a unique value to each allocation algorithm; this value provides an upper and a matching lower bound on the Price of Anarchy of a derived mechanism provable via smoothness. The characterization also applies to the sequential or simultaneous composition of single-parameter mechanisms. Importantly, the factor that we identify is typically not in one-to-one correspondence to the approximation guarantee of the algorithm. Rather, it is usually the product of the approximation guarantee and the degree to which the mechanism is loser independent.

We apply our characterization to show the optimality of greedy mechanisms for single-minded combinatorial auctions, whether these mechanisms are polynomial-time computable or not. We also use it to establish the optimality of a non-greedy, randomized mechanism for independent set in interval graphs and show that it is strictly better than any other deterministic mechanism.

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 studies optimization problems in which the input is held by strategic agents. It can be viewed as algorithm design with the additional twist that the input is provided by selfish agents that need to be incentivized via payments to reveal this information. A natural requirement in this context is truthfulness. The algorithmic requirements for dominant strategy incentive compatibility or Bayes-Nash incentive compatibility are well understood [e.g., Rochet 1987; Myerson 1981; McAfee and McMillan 1988; Jehiel et al. 1996; Jehiel and Moldovanu 2001; Krishna and Maenner 2001; Saks and Yu 2005; Bikhchandani et al. 2006].

Many practical mechanisms, however, are *not* truthful. Examples include the Generalized Second Price (GSP) mechanism for sponsored search auctions (see, e.g., [Edelman et al. 2007; Varian 2007]) or auction protocols for combinatorial auctions that have been used to sell spectrum rights (see, e.g., [Milgrom 2004]). While the exact arguments that have been used to explain the use of non-truthful mechanisms typically depend on the application, there are some general themes. One of these themes

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.

 $Emails: \verb|p.d.duetting@lse.ac.uk| and thomas.kesselheim@mpi-inf.mpg.de.$

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.

 $ACM\ 978\text{-}1\text{-}4503\text{-}3410\text{-}5/15/06\ ...\$15.00.$

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

http://dx.doi.org/10.1145/2764468.2764507

is that these mechanisms are often viewed as *simpler* than their truthful counterparts. For example, a pay-your-bid rule is simpler than a VCG-style payment rule that would be required for truthfulness. Another theme is that the properties that ensure truthfulness are a "global" property of mechanisms that is difficult to achieve. This plays a particular role when one considers that agents typically participate in several mechanisms at once. In this case truthfulness of the individual mechanisms does not necessarily guarantee truthfulness of the global mechanism.

In this paper we do not seek to understand why non-truthful mechanisms are used in practice, but rather which (algorithmic) properties of mechanisms ensure that they perform well despite the strategic behavior of agents. Specifically, we seek to understand which properties of a mechanism ensure that all its equilibria, ideally from a broad range of equilibrium concepts, are close to optimal. That is, we want to understand what keeps the Price of Anarchy of a mechanism with respect to equilibrium concepts such as correlated equilibria or Bayes-Nash equilibria small.

The most versatile technique known to date for proving bounds that apply to these equilibrium concepts is the smoothness framework. Originally formulated by Roughgarden [2012a,b] for games, it was recently generalized to mechanisms by Syrgkanis and Tardos [2013]. Smoothness is a parametrized property, the parameters are $\lambda, \mu \geq 0$, that a mechanism can have and which guarantees that its Price of Anarchy, the worst possible ratio between the optimal welfare and the welfare at equilibrium, is (upper) bounded by $O(\max(1,\mu)/\lambda)$. Another nice aspect about smoothness is that the simultaneous or sequential composition of smooth mechanisms remains smooth [Syrgkanis and Tardos 2013]. What smoothness generally does not achieve is to also show that the bound is best possible. A lower bounding technique, however, is essential for singling out which properties of a mechanism ensure that it achieves near-optimal performance in all equilibria. Without it we can only hope for sufficient conditions, but not for conditions that are both sufficient and necessary.

In fact, most of the previous work gave sufficient conditions. Lucier and Borodin [2010], for example, showed that greedy algorithms with approximation guarantee α have a Price of Anarchy of $O(\alpha)$; while Dütting et al. [2015] showed that this is also true for the algorithm design principle of relax-and-round. In a similar spirit, Babaioff et al. [2014] found that in certain settings declared welfare maximization translates into Price of Anarchy guarantees, and conjectured that this connection could be true much more generally. Specifically, they conjectured that it could be true for approximate declared welfare maximization.

1.1. Our Results

We consider general single-parameter settings where the algorithmic problem consist of choosing a feasible set of "winners" $S \subseteq N$ from a family of feasible sets $\mathcal{F} \subseteq 2^N$ that (approximately) maximizes welfare as given by the sum $\sum_{i \in S} v_i$ of the winners' values v_i for being selected. We focus on monotone algorithms in which stating a higher value can only increase the chance of winning.

These settings are of particular interest for various reasons. First of all, we are not aware of any characterization of algorithms and their Price of Anarchy. Specifically, not even for monotone, single-parameter mechanisms it is understood which algorithmic properties guarantee that all equilibria are close to optimal. A second reason is their potential use in composed mechanisms [Syrgkanis and Tardos 2013]. One instance of such composition is the case of multi-minded bidders. Here, bidders can be served in multiple different ways, which they might value differently. Designing a truthful mechanism for such a setting is far more complex than for the single-parameter setting. In particular, monotonicity in each component is not sufficient for truthfulness. Another instance is when agents participate in several different mechanisms, either

simultaneously or sequentially. In this case designing truthful mechanisms is typically even infeasible as the property that ensures truthfulness is not a "local" property of the individual mechanisms.

Our results will not only apply to individual mechanisms, but will also extend to their composition. Moreover, they will apply for to a broad range of payment schemes, including the pay-your-bid rule in which agents pay their bid if they are chosen and zero otherwise or VCG-style payments in which agents pay their externality. We make the standard assumption that the agents have quasi-linear preferences.

Exact Characterization. Our main result is a tight characterization of single-parameter algorithms, deterministic or randomized, that when paired with suitable payments result in a mechanism with low Price of Anarchy. To this end, we introduce an instance-based measure termed *permeability*. This measure only depends on the feasibility structure and the underlying allocation algorithm, which determines the winners for a given bid vector. It provides an upper and a matching lower bound on the Price of Anarchy provable via smoothness (a.k.a. the robust Price of Anarchy [Roughgarden 2012a]).

In case of a deterministic mechanism, the definition reads as follows. Consider an allocation algorithm $ALG \colon \mathbb{R}^N \to \mathcal{F}$ mapping each vector of bids to a feasible outcome. For a fixed agent $i \in N$ and bids b_{-i} of the agents other than i define $\tau_i(b_{-i})$ as the largest value that this agent can bid and still lose. Then the *permeability* γ of ALG is the smallest factor such that for all feasible sets of agents S,

$$\sum_{i \in S} \tau_i(b_{-i}) \le \gamma \cdot \sum_{i \in ALG(b)} b_i.$$

This and similar properties have been used in previous works to upper-bound the Price of Anarchy by $O(\gamma)$ (see, e.g., [Lucier and Borodin 2010; Syrgkanis and Tardos 2013; Babaioff et al. 2014; Hartline et al. 2014]). Our main technical achievement is to show that the above condition and a suitable generalization to randomized mechanisms yields matching upper and lower bounds. That is, we show that the Price of Anarchy provable via smoothness of a mechanism based on a γ -permeable (randomized) algorithm is $\Theta(\gamma)$. This upper and lower bound also naturally extends to the simultaneous or sequential composition of mechanisms.

In contrast to mechanism smoothness, permeability is a property of the underlying feasibility structure and allocation algorithm. It does not involve the strategic aspects faced by the agents, particularly not the payments and utilities. Showing upper and lower bounds therefore boils down to purely combinatorial arguments, similar to analyzing the approximation ratio of an algorithm.

Optimal Mechanisms. An important feature of our characterization is that by considering a certain problem structure and arguing that the value we identify must be of a specific order for *any* mechanism, one can establish lower bounds that apply to *all* mechanisms.

As a first example of this proof pattern, we provide lower bounds for single-minded combinatorial auctions. The bounds are $\Omega(d)$ and $\Omega(\sqrt{m})$, where d is the maximum bundle size and m is the number of items. These bounds apply to both deterministic and randomized mechanisms. They match the approximation and hence Price of Anarchy guarantees of greedy-by-value and greedy-by-square-root-of-the-bundle-size [Lucier and Borodin 2010]. We thus establish the optimality of greedy algorithms for this problem among all deterministic and randomized mechanisms, whether these algorithms are polynomial-time computable or not.

Then, as a second example, we design a randomized, non-greedy polynomial-time mechanism for problems whose feasibility structure can be expressed as interval

graphs. We use our framework to show that this mechanism has a Price of Anarchy of $O(\log(n))$, where n is the number of players. We prove its optimality by providing an (almost) matching lower bound of $\Omega(\log(n)/\log\log(n))$ that applies to randomized mechanisms. We complement this with a lower bound of $\Omega(\sqrt{n})$ for deterministic mechanisms, which in particular shows that greedy algorithms are provably suboptimal for this problem.

Further Results. We also provide an answer to the question whether (approximate) declared welfare maximization is sufficient. Our characterization implies that generally the Price of Anarchy is *not* in one-to-one correspondence to the approximation guarantee. Rather, as we show, it is typically determined by the product of the approximation guarantee and the degree to which the mechanism is loser independent. We thus not only show what the relevant factors involved in choosing between different algorithms are, but also how the respective factors come together to jointly determine the performance guarantee.

Then, we point out the interesting possibility of indirect algorithmic characterizations of the factor we show uniquely determines an algorithm's applicability in strategic settings. To this end, we show that for all problem structures the corresponding factor for exact declared welfare maximization is given by the approximation guarantee of the greedy-by-value algorithm for this problem. We thus show how our parameter for a certain class of algorithms is uniquely determined by the approximation guarantee of a *different* class of algorithms. On a conceptual level we thus reduce the novel question of understanding the factor that we define to a problem that the theoretical computer science community is used to and well trained in.

Finally, we give some evidence that our characterization also applies to non-smooth mechanisms. Specifically, we show that for deterministic mechanisms, under mild assumptions, there exists a pure Nash equilibrium which achieves the welfare loss that we establish in our upper/lower bounds.

1.2. Related Work

Their is a long body of work that characterizes dominant strategy (or Bayes-Nash) incentive compatibility [e.g., Rochet 1987; Myerson 1981; McAfee and McMillan 1988; Jehiel et al. 1996; Jehiel and Moldovanu 2001; Krishna and Maenner 2001; Saks and Yu 2005; Bikhchandani et al. 2006]. An important difference to our work is that this line of work does not need to worry about multiplicity of equilibria: if truthtelling is an equilibrium this equilibrium provides a focal point for the analysis.

The smoothness framework was introduced by [Roughgarden 2012a,b] for games and extended to mechanisms by Syrgkanis and Tardos [2013]. As of today this technique is the main technique for proving Price of Anarchy bounds that apply to general equilibrium concepts such as correlated equilibria or Bayes-Nash equilibria. Very recently, Kulkarni and Mirrokni [2015] presented a proof pattern based on LP and Fenchel duality that also allows to prove such Price of Anarchy bounds. It is not yet clear how this technique relates to smoothness. As formulated this technique only allows to prove upper bounds, and thus does not provide the lower bounds that would be required for a characterization result.

The closest to our characterization approach are the already mentioned works of Lucier and Borodin [2010] and Dütting et al. [2015] who gave general constructions that show how for larger classes of algorithms the approximation guarantee automatically translates into a Price of Anarchy guarantee that applies to correlated and Bayes-Nash equilibria. The Price of Anarchy of declared welfare maximizers was analyzed in [Dütting et al. 2013] and [Babaioff et al. 2014], while Dütting et al. [2015] characterize

the problem structures and mechanisms that admit a unique correlated equilibrium that achieves optimal social welfare.

Also relevant as a precursor and parallel literature to our work is work that analyzes the Price of Anarchy of simple mechanisms, such as [Christodoulou et al. 2008; Bhawalkar and Roughgarden 2011; Feldman et al. 2013], work that extends the smoothness concept so that it also yields revenue guarantees [Hartline et al. 2014], and work that identifies barriers to near-optimal equilibria [Roughgarden 2014]. The first line of work differs from ours as it considers the performance of simple mechanisms for specific problems. The second line of work suggests an interesting direction for future work, namely extending our characterization to also capture revenue. Finally, the third line of work is in some sense orthogonal to our work as it points out computational rather than information-theoretic barriers.

2. PRELIMINARIES

Mechanism Design Basics. We focus on mechanisms for binary single-parameter problems. Each of n players $i \in N$ holds a single private non-negative number v_i . He can either "win" or "lose"; his valuations for these outcomes are v_i and 0. There is a family of subsets of players $\emptyset \neq \mathcal{F} \subseteq 2^N$, which defines feasible solutions, that is, which players can win simultaneously. In many applications, this set system is downward closed. However, this is not a requirement for our results. The set \mathcal{F} is assumed to be public knowledge. The $social\ welfare$ of a set of winners $S \in \mathcal{F}$ is defined as $\sum_{i \in S} v_i$.

We consider direct mechanisms M, which ask the agents to report their valuations; these reported valuations are referred to as bids, denoted by b. Based on the bids, the mechanism computes a feasible solution $ALG(b) \in \mathcal{F}$. The function ALG is called the allocation algorithm; it may be randomized. We assume it to be monotone in every component. That is, for every player i, every wo bids $b_i \leq b_i'$, and every bid vector b_{-i} , we have $\Pr\left[i \in ALG(b_i, b_{-i})\right] \leq \Pr\left[i \in ALG(b_i', b_{-i})\right]$. Given the bids mechanism M also computes payments $p(b) \in \mathbb{R}^n_{\geq 0}$. A pay-your-

Given the bids mechanism M also computes payments $p(b) \in \mathbb{R}^n_{\geq 0}$. A pay-your-bid mechanism charges winning agents $i \in ALG(b)$ their bid, while losing agents $i \notin ALG(b)$ pay nothing. That is, $p_i(b) = b$ for $i \in ALG(b)$ and $p_i(b) = 0$ otherwise. More generally, we will consider mechanisms whose payment rule p satisfies $0 \le p_i(b) \le b_i$ for $i \in ALG(b)$ and $p_i(b) = 0$ if $i \notin ALG(b)$. That is, mechanisms that do not overcharge the players.

We assume quasi-linear utilities and risk neutrality. For a deterministic mechanism, this means that if bidder i wins under bid vector b, his utility is $v_i - p_i(b)$, otherwise it is 0. More generally $u_i(b) = v_i \mathbf{Pr} \left[i \in ALG(b) \right] - \mathbf{E} \left[p_i(b) \right]$. We study equilibria that are based on these utility functions. Most prominently, this includes pure and mixed Nash equilibria, and their extensions namely Bayes-Nash equilibria and correlated equilibria [Nisan et al. 2007].

Generally, an equilibrium is a probability distribution over bid vectors b. Under incomplete information, the valuations will come from a product distribution $D = D_1 \times \cdots \times D_n$. The complete information setting, where the valuation profile is fixed, corresponds to the special case where the support size of each distribution is one. We write $SW(D,b) = \mathbf{E}_{v \sim D}[\sum_{i \in N} v_i \Pr[i \in ALG(b)]]$ for the expected welfare of a mechanism on input b. We denote the optimal expected welfare by $SW_{opt}(D) = \mathbf{E}_{v \sim D}[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i]$. The $Price\ of\ Anarchy$ with respect to an equilibrium concept $\mathsf{Eq} \in \{\mathsf{PureNash}, \mathsf{MixedNash}, \mathsf{Correlated}, \mathsf{BayesNash}\}$ is the worst possible ratio between the optimal expected welfare and the expected welfare at equilibrium

$$\max_{D,\pi \in \mathsf{Eq}(D)} \frac{SW_{opt}(D)}{\mathbf{E}_{b \sim \pi}[SW(D,b)]}.$$

Smoothness Framework. Syrgkanis and Tardos [2013] defined a mechanism to be weakly (λ, μ_1, μ_2) -smooth for each valuation profile v and every bid profile 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

$$\sum_{i \in N} u_i(b_i', b_{-i}) \ge \lambda \cdot \max_{S \in \mathcal{F}} \sum_{i \in S} v_i - \mu_1 \cdot \sum_{i \in N} \mathbf{E} \left[p_i(b) \right] - \mu_2 \cdot \sum_{i \in N} b_i \mathbf{Pr} \left[i \in f(b) \right] .$$

A mechanism is called (λ, μ) -smooth if it is weakly $(\lambda, \mu, 0)$ -smooth.

Theorem 2.1 (Syrgkanis and Tardos [2013]). If a mechanism is weakly (λ, μ_1, μ_2) -smooth and agents have the possibility to withdraw from the mechanism and do not overbid, then the expected social welfare at any pure or mixed Nash, correlated, or mixed Bayes-Nash equilibrium is at least $\lambda/(\max(\mu_1, 1) + \mu_2)$ of the optimal social welfare. If a mechanism is (λ, μ) smooth, then the overbidding assumption is not required and the bound is $\lambda/\max(\mu, 1)$.

Furthermore, as shown in [Syrgkanis and Tardos 2013], the simultaneous and sequential composition of multiple (weakly) smooth mechanisms remains (weakly) smooth

Following Roughgarden [2012a], we define the *robust Price of Anarchy* of a mechanism M to be the infimum of all ratios $(\max(\mu_1,1) + \mu_2)/\lambda$ such that M is weakly (λ,μ_1,μ_2) -smooth.

3. EXACT CHARACTERIZATION

In this section we give a tight characterization of the Price of Anarchy of a given single-parameter mechanism. The crucial step is the definition of a measure called permeability, which relates the critical bids of groups of players to the reported welfare. Intuitively, it captures how resistant the allocation algorithm is towards losing players. The higher the factor that we define is the less permeable the algorithm.

We then show that permeability is the quintessential property of an allocation algorithm when considering its use in a non-truthful mechanism. It uniquely determines the best Price of Anarchy bound provable via smoothness (a.k.a. the robust Price of Anarchy) for both deterministic and randomized mechanisms.

Finally, we argue that permeability also precisely characterizes the Price of Anarchy provable via smoothness for simultaneous or sequential composition of single-parameter mechanisms.

3.1. Deterministic Mechanisms

As a warm-up and to build up some intuition we first consider deterministic algorithms. For deterministic algorithms $ALG: \mathbb{R}^n \to \mathcal{F}$ we define permeability as follows.

For a fixed agent $i \in N$ and bids b_{-i} of the agents other than i define $\tau_i(b_{-i})$ as the largest value that this agent can bid and still lose. Then the *permeability* γ of ALG is the smallest factor such that for all feasible sets of agents $S \in \mathcal{F}$,

$$\sum_{i \in S} \tau_i(b_{-i}) \le \gamma \cdot \sum_{i \in ALG(b)} b_i.$$

In words: If we take a group of players S and compare the critical bids that these players would have to make to become winners to the reported welfare achieved by the algorithm, then the former can only be larger than the latter by a factor of γ .

Theorem 3.1. A mechanism based on a γ -permeable algorithm is weakly $(\frac{1}{2},0,\gamma)$ -smooth. In particular, a pay-your-bid mechanism based on a γ -permeable algorithm is $(\frac{1}{2},\gamma)$ -smooth.

PROOF. For each bidder i, let $b'_i = \frac{1}{2}v_i$. Furthermore, let S = OPT(v). Observe that for some $i \in S$, we always have $u_i(b'_i, b_{-i}) \geq \frac{1}{2}v_i - \tau_i(b_{-i})$, no matter if it is selected by the algorithm or not. Summing over all bidders in S, we get

$$\sum_{i \in N} u_i(b'_i, b_{-i}) \ge \sum_{i \in S} u_i(b'_i, b_{-i}) \ge \frac{1}{2} \sum_{i \in S} v_i - \sum_{i \in S} \tau_i(b_{-i}) \ge \frac{1}{2} \sum_{i \in S} v_i - \gamma \sum_{i \in ALG(b)} b_i . \quad \Box$$

We thus obtain an upper bound of $O(\gamma)$ on the Price of Anarchy of a mechanism based on a γ -permeable algorithm. Surprisingly, the permeability γ of an algorithm can also be used to prove a matching lower bound of $\Omega(\gamma)$. Before we prove this let us first see how to generalize the definition of permeability to randomized mechanisms.

3.2. Randomized Mechanisms

In contrast to a deterministic algorithm, a randomized algorithm usually does not have clear breakpoints such as the critical values $\tau_i(b_i)$. Instead, the winning probability might increase stepwise or even smoothly. To take this into consideration, we define a critical value for each probability $0 \le q \le 1$ as follows. Given a bid vector b, for some bidder $i \in N$ and $0 \le q \le 1$, let $\tau_i(b_{-i},q) = \sup\{b_i' \mid \mathbf{Pr} \left[i \in ALG(b_i',b_{-i})\right] < q\}$. That is, roughly spoken, $\tau_i(b_{-i},q)$ is the highest value that i can bid against b such that the winning probability remains below q.

The *permeability* of a monotone randomized algorithm is now the infimum over all numbers $\gamma \geq 1$ such that for any bid vector b and any feasible set S

$$\sum_{i \in S} \tau_i \left(b_{-i}, \frac{1}{\gamma} \right) \le \gamma \sum_{i \in N} b_i \mathbf{Pr} \left[i \in ALG(b) \right] .$$

Observe that larger values of γ loosen this condition in two ways: Larger values of γ cause each $\tau_i(b_{-i},\frac{1}{\gamma})$ term and hence the sum on the left-hand side to be smaller. At

the same time, the right-hand side increases with γ . However, for a deterministic algorithm $\tau_i\left(b_{-i},\frac{1}{\gamma}\right)=\tau_i(b_{-i})$ for any $\gamma\geq 1$. Therefore, this definition of permeability is consistent with the one stated for deterministic algorithms.

First we show that using a γ -permeable (randomized) algorithm is sufficient to show a bound of $O(\gamma)$ on the Price of Anarchy.

THEOREM 3.2. A mechanism based on a γ -permeable algorithm is weakly $(\frac{1}{2\alpha}, 0, 1)$ smooth. In particular, a pay-your-bid mechanism based on a γ -permeable algorithm is $(\frac{1}{2\gamma},1)$ -smooth.

PROOF. For each bidder i, let $b'_i = \frac{1}{2}v_i$. Furthermore, let S = OPT(v). If $\mathbf{Pr}\left[i \in ALG(b_i',b_{-i})\right] < \frac{1}{\gamma}$, then $\tau_i\left(b_{-i},\frac{1}{\gamma}\right) \geq b_i'$. Thus $u_i(b_i',b_{-i}) \geq 0 \geq$ $\frac{1}{\gamma} \left(\frac{1}{2} v_i - \tau_i \left(b_{-i}, \frac{1}{\gamma} \right) \right).$ If $\mathbf{Pr}\left[i \in ALG(b'_i, b_{-i})\right] \geq \frac{1}{\gamma}$, then

$$u_i(b'_i, b_{-i}) = \frac{1}{2} v_i \mathbf{Pr} \left[i \in ALG(b'_i, b_{-i}) \right] \ge \frac{1}{2} v_i \frac{1}{\gamma} \ge \frac{1}{\gamma} \left(\frac{1}{2} v_i - \tau_i \left(b_{-i}, \frac{1}{\gamma} \right) \right) .$$

Summing over all bidders in *S*, we get

$$\begin{split} \sum_{i \in N} u_i(b_i', b_{-i}) &\geq \sum_{i \in S} u_i(b_i', b_{-i}) \\ &\geq \frac{1}{2\gamma} \sum_{i \in S} v_i - \sum_{i \in S} \frac{1}{\gamma} \tau_i \left(b_{-i}, \frac{1}{\gamma} \right) \\ &\geq \frac{1}{2\gamma} \sum_{i \in S} v_i - \sum_{i \in N} b_i \mathbf{Pr} \left[i \in ALG(b) \right] . \quad \Box \end{split}$$

Interestingly, this bound is tight on a per-instance basis: Any smoothness-based proof will always imply a guarantee of $\Omega(\gamma)$ on the Price of Anarchy.

THEOREM 3.3. A mechanism based on a γ -permeable algorithm is not weakly (λ, μ_1, μ_2) -smooth for any λ , μ_1 , and μ_2 such that $\frac{\lambda}{2(\mu_1 + \mu_2 + 1)} > \frac{1}{\gamma}$.

PROOF. Given any b and S making the bound tight. For some arbitrary $\epsilon>0$, set $v_i=\tau_i\left(b_{-i},\frac{1}{\gamma}\right)-\epsilon$ for $i\in S$, and $v_i=b_i$ otherwise.

Let b' be an arbitrary bid vector. We assume that $b'_i > v_i$ is weakly dominated. Therefore, monotonicity gives us $\mathbf{E}[u_i(b'_i,b_{-i})] \leq v_i \mathbf{Pr}[i \in ALG(v_i,b_{-i})]$.

For $i \in S$, this implies $\mathbf{E}\left[u_i(b_i',b_{-i})\right] \leq v_i\mathbf{Pr}\left[i \in ALG(v_i,b_{-i})\right] \leq \frac{1}{\gamma}\tau_i\left(b_{-i},\frac{1}{\gamma}\right)$. For $i \notin S$, we naturally have $\mathbf{E}\left[u_i(b_i',b_{-i})\right] \leq v_i\mathbf{Pr}\left[i \in ALG(v_i,b_{-i})\right] = b_i\mathbf{Pr}\left[i \in ALG(b)\right]$. As b and S are chosen to make the bound tight, we have $\sum_{i \in N} b_i\mathbf{Pr}\left[i \in ALG(b)\right] = \frac{1}{\gamma}\sum_{i \in S}\tau_i\left(b_{-i},\frac{1}{\gamma}\right)$. This implies for all $\mu_1,\mu_2 \geq 0$

$$\begin{split} &\sum_{i \in N} \mathbf{E} \left[u_{i}(b'_{i}, b_{-i}) \right] + (\mu_{1} + \mu_{2}) \sum_{i \in N} b_{i} \mathbf{Pr} \left[i \in ALG(b) \right] \\ &\leq \sum_{i \in S} \frac{1}{\gamma} \tau_{i} \left(b_{-i}, \frac{1}{\gamma} \right) + (\mu_{1} + \mu_{2} + 1) \sum_{i \in N} b_{i} \mathbf{Pr} \left[i \in ALG(b) \right] \\ &\leq (\mu_{1} + \mu_{2} + 1) \left(\sum_{i \in S} \frac{1}{\gamma} \tau_{i} \left(b_{-i}, \frac{1}{\gamma} \right) + \sum_{i \in N} b_{i} \mathbf{Pr} \left[i \in ALG(b'_{i}, b_{-i}) \right] \right) \\ &= \frac{2(\mu_{1} + \mu_{2} + 1)}{\gamma} \sum_{i \in S} \tau_{i} \left(b_{-i}, \frac{1}{\gamma} \right) \\ &= \frac{2(\mu_{1} + \mu_{2} + 1)}{\gamma} \sum_{i \in S} v_{i} + \frac{2(\mu_{1} + \mu_{2} + 1) |S| \epsilon}{\gamma} \\ &\leq \frac{2(\mu_{1} + \mu_{2} + 1)}{\gamma} \sum_{i \in OPT(v)} v_{i} + \frac{2(\mu_{1} + \mu_{2} + 1) |S| \epsilon}{\gamma} \end{split} .$$

On the other hand, to make the mechanism weakly (λ, μ_1, μ_2) -smooth, we need

$$\sum_{i \in N} \mathbf{E} [u_i(b'_i, b_{-i})] \ge \lambda \sum_{i \in OPT(v)} v_i - \mu_1 \sum_{i \in N} \mathbf{E} [p_i(b)] - \mu_2 \sum_{i \in N} b_i \mathbf{Pr} [i \in ALG(b'_i, b_{-i})]$$

$$\ge \lambda \sum_{i \in OPT(v)} v_i - (\mu_1 + \mu_2) \sum_{i \in N} b_i \mathbf{Pr} [i \in ALG(b'_i, b_{-i})] .$$

In combination, we get

$$\lambda \sum_{i \in OPT(v)} v_i \le 2(\mu_1 + \mu_2 + 1) \frac{1}{\gamma} \sum_{i \in OPT(v)} v_i + \frac{2(\mu_1 + \mu_2 + 1)|S|\epsilon}{\gamma}.$$

As this bound holds for all $\epsilon>0$, we also have $\lambda\leq \frac{2(\mu_1+\mu_2+1)}{\gamma}$. $\ \ \, \Box$

3.3. Composition of Mechanisms

Syrgkanis and Tardos [2013] have shown that the simultaneous or sequential composition of mechanisms remains smooth. Applied to single parameter mechanisms there result for the simultaneous composition of mechanisms says that whenever agents have XOS, or fractionally subadditive valuations, then the simultaneous composition of k mechanisms M_j that are $(\lambda^{(j)},\mu^{(j)})$ -smooth is $(\min_j \lambda^{(j)},\max_j \mu^{(j)})$ -smooth. Similarly, applied to the sequential composition of single-parameter mechanisms, there result says that if agents have unit demand valuations, then the sequential composition of k single-parameter mechanisms M_j that are $(\lambda^{(j)},\mu^{(j)})$ -smooth leads to a mechanism that is $(\min_j \lambda^{(j)},\max_j \mu^{(j)}+1)$ -smooth.

Since, as we have shown, every pay-your-bid mechanism based on a γ -permeable mechanism is $(1/2\gamma,1)$ -smooth, we can user their result to show that the simultaneous or sequential composition of k pay-your-bid mechanisms based on γ_1,\ldots,γ_k -permeable algorithms is $(1/(2\max_j\gamma_j),1)$ -smooth resp. $(1/(2\max_j\gamma_j),2)$ -smooth. Maybe more surprisingly, not only our upper bound but also our lower bound carries over to compositions. For this we simply apply the lower bound proof to the pay-your-bid mechanism that is based on the algorithm with the worst γ_j .

Both the results of Syrgkanis and Tardos and the above arguments that show that our upper and lower bounds carry over to compositions apply equally well to weakly smooth mechanisms.

4. OPTIMAL MECHANISMS

Due to the characterization via permeability, we can bound the Price of Anarchy of mechanisms by only arguing about the allocation algorithm, abstracting from the strategic aspects such as private valuations and bids. This is particularly helpful to derive impossibility results. In this section, we demonstrate this with two examples. First, we consider the well-known setting of combinatorial auctions. Here, we can show that existing greedy algorithms perform optimally in terms of permeability, even compared to unlimited computational power. That is, these algorithms achieve the best-possible robust Price of Anarchy.

We also study the weighted independent-set problem in interval graphs. This problem is solvable in polynomial time. However, no deterministic algorithm is γ -permeable for $\gamma \in o(\sqrt{n})$. With randomization, there is an $O(\log n)$ -permeable algorithm. We will show that this is almost optimal, too.

4.1. Combinatorial Auctions

In a combinatorial auction, m items are sold. Each bidder i is interested a (publicly known) bundle of items $T_i \subseteq [m]$. A set of bidders $S \subseteq N$ is feasible (i.e., $S \in \mathcal{F}$) if and only if $T_i \cap T_{i'} = \emptyset$ for all $i, i' \in S$, $i \neq i'$. By d we denote the maximum bundle size $\max_{i \in N} |T_i|$.

Simple greedy algorithms achieve approximation factors $O(\sqrt{m})$ for the case of m items and O(d) [Lehmann et al. 2002]. The results by Lucier and Borodin [2010] and Syrgkanis and Tardos [2013] imply that these bounds transfer to robust-price-of-anarchy guarantees.

THEOREM 4.1. There is no algorithm for combinatorial auctions that is γ -permeable for $\gamma = o(\sqrt{m})$ or $\gamma = o(d)$.

PROOF. To show the statement, let without loss of generality d be a prime number and set $m=d^2$. We show that there is an instance with d^2 bidders and d^2 items such that for every algorithm there is a bid vector b for which

$$\sum_{i \in S} \tau_i \left(b_{-i}, \frac{2}{d} \right) \ge \frac{d}{2} \mathbf{E} \left[\sum_{i \in N} b_i(f(b)) \right]$$

for some feasible set S.

We consider the finite field \mathbb{F}_d . We identify both the set of items and the set of bidders by $\mathbb{F}_d \times \mathbb{F}_d$. Bidder $(i,j) \in \mathbb{F}_d \times \mathbb{F}_d$ is interested in buying set $T_{(i,j)} = \{(x,ix+j) \mid x \in \mathbb{F}_d\}$. Observe that for $(i,j) \neq (i',j')$, we have $T_{(i,j)} \cap T_{(i',j')} = \emptyset$ if and only if i=i'. That is, feasible solutions are precisely subsets of bidders that share the same first index. Furthermore, for some $z \in \mathbb{F}_d^d$, let us further define a bid vector b^z by setting $b_{(i,j)}^z = 1$

Furthermore, for some $z \in \mathbb{F}_d^d$, let us further define a bid vector b^z by setting $b_{(i,j)}^z = 1$ if $z_i = j$ and 0 otherwise. This means that in b^z no two bidders bidding 1 share the first same index. Therefore any pair of bidders bidding 1 simultaneously exclude each other.

To show the existence of the claimed b, we use the probabilistic method. We draw z uniformly at random from \mathbb{F}_d^d and i uniformly at random from \mathbb{F}_d independently. In the bid vector b^z , the winning probabilities of bidders bidding 1 sum up to at most 1 due to mutual exclusion. Therefore, the expected winning probability of (i, z_i) is at most $\frac{1}{d}$.

Next, we describe an alternative way of drawing z. For this purpose, draw \tilde{z} uniformly at random from \mathbb{F}_d^d . Again, i uniformly at random from \mathbb{F}_d independently. Furthermore, j is now drawn independently uniformly from \mathbb{F}_d . We derive z by setting $z_i = j$ and $z_{i'} = \tilde{z}_{i'}$ for $i' \neq i$.

We already observed that the expected winning probability of (i, z_i) in b^z is at most $\frac{1}{d}$. This means, there have to be \hat{z} and \hat{i} such that, conditioned on $\tilde{z} = \hat{z}$ and $i = \hat{i}$, the expected winning probability of (\hat{i}, j) in b^z is at most $\frac{1}{d}$.

We now define b by setting $b_{(i',j')}=1$ for $i'\neq\hat{i}$ and $\tilde{z}_{i'}=j'$ and 0 otherwise. By the above construction, the expected winning probability of (\hat{i},j) in $(1,b_{-(\hat{i},j)})$ is at most $\frac{1}{d}$. By Markov's inequality this means that with probability at least $\frac{1}{2}$ it is less than $\frac{2}{d}$. In such a case, $\tau_{(\hat{i},j)}(b_{-(\hat{i},j)},\frac{2}{d})\geq 1$. That is,

$$\sum_{j' \in \mathbb{F}_d} \tau_{(\hat{i},j')} \left(b_{-(\hat{i},j')}, \frac{2}{d} \right) \ge \frac{d}{2} . \quad \Box$$

4.2. Independent Set in Interval Graphs

We now turn to the independent-set problem in interval graphs, which can also be considered a more structured version of combinatorial auctions. Formulated as a single-parameter mechanism design problem, the task is as follows. Each bidder i is interested in a (publicly known) interval of real numbers $T_i = [l_i, r_i] \subseteq \mathbb{R}$. Again, $S \in \mathcal{F}$ if and only if $T_i \cap T_{i'} = \emptyset$ for all $i, i' \in S$, $i \neq i'$. So, this problem corresponds to a combinatorial auction where all bidders desire sets T_i of connected items.

THEOREM 4.2. For independent set in interval graphs, there is no γ -permeable deterministic algorithm for $\gamma = o(\sqrt{n})$.

PROOF. Let $T_i=[i,i]$ for $1\leq i\leq n-1$ and $T_n=[1,n-1]$. Consider a fixed γ -permeable algorithm. Observe that its declared welfare on $b=(0,\dots,0,1)$ is at most 1. Therefore, we have $\sum_{i=1}^{n-1} \tau_i(b_{-i}) \leq \gamma$.

Now consider some $i \leq n-1$ and $\epsilon > 0$ and the bid vector $(\tau_i(b_{-i}) + \epsilon, b_{-i})$. Under this bid vector bidder i wins. That is, the declared welfare achieved by the algorithm is $\tau_i(b_{-i}) + \epsilon$. Bidder n is now a loser. By monotonicity of the algorithm his critical value under the modified bid vector has to be at least 1. This means, we have $1 \leq \gamma(\tau_i(b_{-i}) + \epsilon)$ for all $i \leq n-1$ and thus $n-1 \leq \gamma \sum_{i=1}^{n-1} (\tau_i(b_{-i}) + \epsilon)$.

for all $i \leq n-1$ and thus $n-1 \leq \gamma \sum_{i=1}^{n-1} (\tau_i(b_{-i}) + \epsilon)$. Multiplying both obtained inequalities, we get $(n-1) \sum_{i=1}^{n-1} \tau_i(b_{-i}) \leq \gamma^2 \sum_{i=1}^{n-1} (\tau_i(b_{-i}) + \epsilon)$ for all $\epsilon > 0$. By continuity, this implies $(n-1) \sum_{i=1}^{n-1} \tau_i(b_{-i}) \leq \gamma^2 \sum_{i=1}^{n-1} \tau_i(b_{-i})$ and thus $n-1 \leq \gamma^2$. \square

Indeed this proof directly breaks if one allows randomization. This is due to the fact that there are only two Pareto-optimal solutions. Randomizing between the two of them would give a 2-permeable algorithm. Indeed this idea generalizes as follows.

THEOREM 4.3. For independent set in interval graphs, there is an $O(\log n)$ -permeable randomized algorithm.

PROOF. Given a bid vector b, the algorithm uses the maximum bid $b_{\max} = \max_{i \in N} b_i$ to divide bidders into classes. For $j \in \mathbb{N}$, let $U_j = \{i \in N \mid b_i \geq 2^{-j+1}b_{\max}\}$ be the set of bidders whose bid is at least $2^{-j}b_{\max}$. The algorithm proceeds by first applying the greedy algorithm that finds the maximum unweighted independent set on each set U_j for $j \in [k]$ where $k = 2\lceil \log_2 n \rceil + 1$. Let the respective outcome be denoted by $G_j(b)$. Next, it defines $H_j(b) = G_j(b) \setminus \bigcup_{j' < j} H_{j'}(b)$. Finally, it draws J uniformly at random from [k] and returns $H_J(b)$. This algorithm is monotone because the probability of being selected is either 0 or $\frac{1}{k}$ and increasing the bid can never reduce the probability.

We claim that this algorithm is $O(\log n)$ -permeable. That is, we have to show that for any set $S \in \mathcal{F}$ and any bid vector b, we have

$$\sum_{i \in S} \tau_i \left(b_{-i}, \frac{1}{k} \right) = O(\log n) \sum_{i \in N} b_i \mathbf{Pr} \left[i \in ALG(b) \right] .$$

Observe that $\tau_i\left(b_{-i},\frac{1}{k}\right) \leq 2b_{\max}$. This is due to the fact that for any $\epsilon > 0$, bidder i has a separate class on $(2b_{\max} + \epsilon, b_{-i})$ and is thus selected with probability $\frac{1}{k}$.

Let $S'=\{i\in S\mid \mathbf{Pr}\ [i\in ALG(b)]<\frac{1}{k}\}$. For some $\epsilon>0$, define a modified bid vector b' by setting $b'_i=\max\{\frac{1}{2}\tau_i\left(b_{-i},\frac{1}{k}\right)-\epsilon,b_i\}$ otherwise. By definition $\mathbf{Pr}\ [i\in ALG(b'_i,b_{-i})]<\frac{1}{k}$ for all $i\in S'$. Furthermore, as $b'_i\leq b_{\max}$, under all bid vectors (b'_i,b_{-i}) , the class structure remains unchanged. This means that some bidder $i\in S'$ can be included in more U_j sets in (b'_i,b_{-i}) than in b but in none of them he is selected by the greedy algorithm for the unweighted problem. Adding bidder that are not selected does not affect the way a greedy algorithm computes its solution. Therefore, the same effect occurs if all $i\in S'$ move simultaneously to b'_i . Formally, this means that ALG(b) and ALG(b') are identically distributed, in particular $\mathbf{Pr}\ [i\in ALG(b')]<\frac{1}{k}$ for all $i\in S'$.

We now have

$$\sum_{i \in S} \tau_i \left(b_{-i}, \frac{1}{k} \right) \le \sum_{i \in S} (2b_i' + \epsilon) .$$

Furthermore, the algorithm is an $O(\log n)$ -approximation for the following reason. On the one hand, the optimum is bounded by always taking the respective upper splitting points for the classes and $U_i \cap OPT(b')$ is always a feasible solution to the

respective unweighted independent set problem, i.e., $|G_j(b')| \ge |U_j \cap OPT(b')|$.

$$\sum_{i \in OPT(b')} b'_i \leq \sum_{j=1}^k 2^{-j+2} b_{\max} |U_j \cap OPT(b')| + \sum_{i \notin U_k} b'_i$$

$$\leq \sum_{j=1}^k 2^{-j+2} b_{\max} |U_j \cap OPT(b')| + 2^{-k} b_{\max} n$$

$$\leq \sum_{j=1}^k 2^{-j+2} b_{\max} |G_j(b')| + 2^{-k} b_{\max} n$$

$$\leq 4k \mathbf{E} \left[2^J b_{\max} |G_J(b')| \right] + \frac{b_{\max}}{n} .$$

On the other hand, the geometric series gives us

$$b_i' \mathbf{Pr} [i \in ALG(b')] = \frac{1}{k} b_i' \left| \bigcup_j H_j(b') \cap \{i\} \right| \ge \frac{1}{k} \sum_{j=1}^k 2^{-j} b_{\max} |G_j(b') \cap \{i\}| ,$$

and therefore

$$\sum_{i \in N} b_i' \mathbf{Pr} \left[i \in ALG(b') \right] \ge \mathbf{E} \left[2^{-J} b_{\max} |G_J(b')| \right] .$$

As trivially $\sum_{i \in OPT(b')} b'_i \geq b_{\max}$, we get

$$\sum_{i \in N} b_i' \mathbf{Pr} \left[i \in ALG(b') \right] \ge \frac{1}{4k(1 - \frac{1}{n})} \sum_{i \in OPT(b')} b_i' .$$

This implies

$$\sum_{i \in N} b_i \mathbf{Pr} \left[i \in ALG(b) \right] = \sum_{i \in N} b_i' \mathbf{Pr} \left[i \in ALG(b') \right] = \Omega \left(\frac{1}{\log n} \right) \sum_{i \in OPT(b')} b_i' \ge \Omega \left(\frac{1}{\log n} \right) \sum_{i \in S} b_i' \ .$$

So, in combination

$$\sum_{i \in S} \tau_i \left(b_{-i}, \frac{1}{k} \right) = O(\log n) \sum_{i \in N} b_i \mathbf{Pr} \left[i \in ALG(b) \right] + n\epsilon .$$

By taking the limit for ϵ to 0, this yields the claim. \Box

However, the $O(\log n)$ -bound is also almost optimal.

Theorem 4.4. For independent set in interval graphs, there is no γ -permeable randomized algorithm for $\gamma = o\left(\frac{\log n}{\log \log n}\right)$.

PROOF. We will show that there is an instance such that for each algorithm there is a bid vector b for which

$$\sum_{i \in S} \tau_i(b_{-i}, q) = \Omega(\log n) \mathbf{E} \left[\sum_{i \in N} b_i(f(b)) \right]$$

for some feasible set S and $q = O\left(\frac{\log \log n}{\log n}\right)$.

The instance is a interval graph of nested intervals. It is constructed recursively as follows: We start on level 0 with only a single interval. Level $\ell+1$ is derived by placing $r=\log^4 n$ intervals next to each other into each interval of level ℓ . This way, we obtain a variant of an r-ary tree with shortcuts. Its height is $h=\Theta(\frac{\log n}{\log r})=\Theta(\frac{\log n}{\log\log n})$.

To show the existence of b for an arbitrary algorithm, we use the probabilistic method. To this end, we will mark a subset $T \subseteq N$ of intervals. For each such marking T, an associated bid vector p^T is defined as follows. For a bidder i corresponding to an interval on level ℓ , we set $b_i^T = r^{-\ell}$ if $i \in T$ and $b_i^T = 0$.

Let us first consider the bid vector $b^{\tilde{T}}$ induced by a marking \tilde{T} obtained as follows. On level $\ell>0$, mark $r^{\ell-1/2}$ intervals chosen uniformly at random out of all r^{ℓ} intervals on this level.

Let \mathcal{E} be the event that in every level $\ell-1$ each interval contains a marked interval of level ℓ . This event occurs with high probability because the probability that a fixed interval in level $\ell-1$ does not contain a marked interval of level ℓ is at most

$$\left(1 - \frac{1}{r^{\ell - 1}}\right)^{r^{\ell - 1/2}} = \left(\left(1 - \frac{1}{r^{\ell - 1}}\right)^{r^{\ell - 1}}\right)^{\sqrt{r}} \ge e^{-\sqrt{r}} \le \frac{1}{n^2}.$$

Fix a set \tilde{T} that fulfills \mathcal{E} . Observe that every marked interval belongs to an h-clique of marked intervals. That is, running the algorithm on $b^{\tilde{T}}$, the average probability of some $i \in \tilde{T}$ being picked at most $\frac{1}{h}$. Therefore, randomizing over all possible vectors \tilde{T} respectively $b^{\tilde{T}}$, the average probability of being picked is at most $\frac{1}{h} + \mathbf{Pr}\left[\bar{\mathcal{E}}\right] \leq \frac{2}{h}$.

This means, there has to be a partial marking T of the following kind. The marking is complete on all levels except one level ℓ . In level ℓ all markings but one are present, that is, $r^{\ell-1/2}-1$ are marked. Now, choosing uniformly one of the remaining intervals on level ℓ to be marked (and completing \tilde{T}), the expected probability of being selected by the algorithm is at most $\frac{2}{h}$.

Let $b = b^T$ be the bid vector induced by this partial marking. The complete marking \tilde{T} is obtained by choosing one additional interval from level ℓ uniformly. Let i be the random variable indicating the index of this interval, and let P_i be its probability to be random variable indicating the index of this interval, and let P_i be its probability to be selected by the algorithm. We now have $\mathbf{E}\left[P_i\right] \leq \frac{2}{h}$. That is, by Markov's inequality, we have $\mathbf{Pr}\left[P_i \geq \frac{4}{h}\right] \leq \frac{1}{2}$ and thus $\mathbf{Pr}\left[P_i < \frac{4}{h}\right] \geq \frac{1}{2}$. Observe that by monotonicity $P_i < \frac{4}{h}$ implies $\tau_i(b_{-i}, \frac{4}{h}) \geq r^{-\ell}$. Therefore, we have $\mathbf{E}\left[\tau_i(b_{-i}, \frac{4}{h})\right] \geq \frac{1}{2}r^{-\ell}$. Let S be the set of possible last indices chosen. The bound on $\mathbf{E}\left[\tau_i(b_{-i}, \frac{4}{h})\right]$ implies $\sum_{i \in S} \tau_i(b_{-i}, \frac{4}{h}) \geq \frac{1}{2}|S|r^{-\ell}$. Furthermore, $|S| = r^{\ell} - r^{\ell-1/2} + 1 \geq \frac{1}{2}r^{\ell}$ for sufficiently large r. This implies $\sum_{i \in S} \tau_i(b_{-i}, \frac{4}{h}) \geq \frac{1}{4}$. The declared welfare of any feasible solution is bounded by $\sum_{\ell=1}^h r^{\ell-1/2} r^{-\ell} = \frac{h}{\sqrt{r}}$.

That is, $\frac{1}{4} \leq \gamma \frac{h}{\sqrt{r}}$ and thus $\gamma \geq \frac{\sqrt{r}}{4h} = \Omega(\log n)$. \square

5. LOSER INDEPENDENCE AND APPROXIMATION GUARANTEE

Using the concept of permeability also allows us to gain a deeper understanding why greedy algorithms perform extremely well when used as mechanisms, as shown by Lucier and Borodin [2010]. The underlying property of greedy algorithms is loser independence. Informally spoken this means that a group of losers stays losers, even if they increase their bids simultaneously while staying under the individual critical value.

We parameterize this property as follows. For a given bid vector b and every set of losing players $T \not\subseteq ALG(b)$ define $\tau_T(b_{-T}) = \sup\{\sum_{i \in T} b_i' \mid ALG(b) = ALG(b_T', b_{-T})\}$. That is, is $\tau_T(b_{-T})$ is the largest sum of bids $\sum_{i \in T} b_i'$ that the players in T could submit without affecting the outcome of the algorithm. An algorithm ALG is β -loser independent if for all sets $T \not\subseteq ALG(b)$

$$\tau_T(b_{-T}) \ge \frac{1}{\beta} \sum_{i \in T} \tau_i(b_{-i}).$$

Note that for algorithms ALG such that $ALG(b) \neq N$ for at least one b we must have $\beta \geq 1$, since for $T = \{i\} \not\subseteq ALG(b)$, $\tau_T(b_{-T}) = \tau_i(b_{-i})$ and hence the previous inequality would be violated if $\beta < 1$. For consistency, in the pathological case that ALG(b) = N for all b, we set $\beta = 1$. In the full version, we show the following connection between permeability, loser independence, and the approximation factor.

THEOREM 5.1. Consider an α -approximation algorithm that is β -loser independent. Then the algorithm is γ -permeable with $\gamma \leq \alpha \beta$.

We also show that if α and β are tight on the same instance, then we get a "reverse" of the previous theorem showing that $\gamma \geq \alpha \beta$.

PROPOSITION 5.2. Consider an α -approximation algorithm ALG that is β -loser independent. If there exists a vector of bids b such that for a set of players $T \not\subseteq ALG(b)$ and the bids b'_i of these players for which $\tau_{-T}(b_{-T}) = \sum_{i \in T} \tau_i(b_{-i})$,

$$au_T(b_{-T}) = rac{1}{eta} \sum_{i \in T} au_i(b_{-i}) \; \; extit{and} \; \; \sum_{i \in ALG(b)} b_i = rac{1}{lpha} \sum_{i \in T} b_i'.$$

then the algorithm is γ -permeable with $\gamma \geq \alpha \beta$.

Two examples for which α and β are tight on the same instance can be found in the context of single-minded combinatorial auctions.

The first example is the algorithm that maximizes declared welfare. The instance is "1-vs-k" with m small bidders and 1 big bidder. The values are one for all players. The big bidder bids 1 and the small bidders bid 0. The set T comprises all small bidders. For each $i \in T$ we have $\tau_i(b_{-i}) = 1$, while $\tau_T(b_{-T}) = 1$. This gives a β -value of m. The α -value is 1.

The second example is the greedy algorithm which ranks players by bid divided through the square root of the bundle size. As before the instance is "1-vs-k" with 1 big bidder and m small bidders. The values are \sqrt{m} and 1. The bids are \sqrt{m} and 0. The set T again comprises all small bidders. Then, for each $i \in T$ we have $\tau_i(b_{-i}) = 1$, while $\tau_T(b_{-T}) = m$. This gives an α -value of \sqrt{m} . The β -value is 1.

6. COMPARISON OF EXACT OPTIMIZATION TO GREEDY-BY-VALUE

Another interesting aspect is that for exact declared welfare maximization permeability can often be analyzed through the approximation guarantee of the greedy algorithm that ranks players by value, from high to low, and accepts the next player if it is feasible to do so. The connection is that, up to constant factors, the approximation guarantee of greedy-by-value provides an upper bound on the permeability γ of exact optimization.

Theorem 6.1. If exact optimization over a downward closed set system \mathcal{F} is γ -permeable, then greedy-by-value is a $\gamma+1$ -approximation.

PROOF. For a bid vector b, let OPT(b) be the optimal solution and ALG(b) be the greedy-by-value solution. Without loss of generality, let $b_1 \geq b_2 \geq \ldots$ with ties broken the same way as the greedy-by-value algorithm does. Define \tilde{b} by $\tilde{b}_i = b_i$ if $i \in ALG(b)$ and 0 otherwise.

Let us consider some $i \not\in ALG(b)$ and $\epsilon > 0$. We claim that $i \not\in OPT(b_i - \epsilon, \tilde{b}_{-i})$. To this end, we will assume the contrary and show that $OPT(b_i - \epsilon, \tilde{b}_{-i})$ would not be optimal under these circumstances. By definition of the greedy algorithm, $\{j \in ALG(b) \mid j < i\} \cup \{i\}$ is infeasible. Therefore, there has to be some $k \in ALG(b)$, k < i, such that $k \not\in OPT(b_i - \epsilon, \tilde{b}_{-i})$. By definition $b_k > b_i - \epsilon$. That is,

$$b_i - \epsilon + \sum_{\substack{j \in OPT(b_i - \epsilon, \tilde{b}_{-i}) \\ j \neq i}} \tilde{b}_j < b_k + \sum_{\substack{j \in OPT(b_i - \epsilon, \tilde{b}_{-i}) \\ j \neq i}} \tilde{b}_j \leq b_k + \sum_{\substack{j \neq k}} \tilde{b}_j = \sum_{\substack{j \in ALG(b)}} \tilde{b}_j \ .$$

This is a contradiction.

As a consequence, we get for any $i \notin ALG(b)$ that $\tau_i(\tilde{b}_{-i}) \geq b_i - \epsilon$ for all $\epsilon > 0$ and therefore $\tau_i(\tilde{b}_{-i}) \geq b_i$. That is, we get

$$\sum_{i \in OPT(b) \backslash ALG(b)} b_i \leq \sum_{i \in OPT(b) \backslash ALG(b)} \tau_i(\tilde{b}_{-i}) \leq \gamma \sum_{i \in OPT(\tilde{b})} \tilde{b}_i = \gamma \sum_{i \in ALG(b)} b_i \enspace,$$

and therefore

$$\sum_{i \in OPT(b)} b_i \le (\gamma + 1) \sum_{i \in ALG(b)} b_i . \quad \Box$$

7. EXTENSION BEYOND SMOOTHNESS

We have shown how our characterization settles the Price of Anarchy provable via smoothness through matching upper and lower bounds. Next we will see how the same insights can be used to tackle the question of whether or not there is a gap between the best bound provable via smoothness and the best bound provable via any technique.

In fact, as we show in the full version, for deterministic mechanisms, under very mild assumptions, there is no gap. Our first result is for mechanisms that charge winning agents their "critical value", i.e., the smallest bid $\tau_i(b_{-i})$ that they would have to win against bids b_{-i} .

PROPOSITION 7.1. For any deterministic mechanism based on a γ -permeable algorithm that uses critical-value payments there is a pure Nash equilibrium that extracts at most a $1/\gamma$ fraction of the optimal social welfare.

Our second result concerns pay-your-bid mechanisms that are defined for a class of instances \mathcal{I} , where each instance $I \in \mathcal{I}$ consists of a set of players N and a feasibility structure \mathcal{F} , and that satisfy a mild consistency assumption.

We assume that any instance $I=(N,\mathcal{F})\in\mathcal{I}$ can be extended to an instance $I'=(N',F')\in\mathcal{I}$ as follows: (1) for each player $i\in N$ add i and a copy i' to N', and (2) augment the feasibility structure so that player i and its copy i' are mutually exclusive, i.e., $i,i'\in S\Rightarrow S\not\in\mathcal{F}'$, but otherwise play identical roles with respect to the feasibility structure, i.e., for any set $S\subseteq N'\setminus\{i\}, S\cup\{i\}\in\mathcal{F}'\Leftrightarrow S\cup\{i'\}\in\mathcal{F}'$ and vice versa. We refer to I,I' as the original instance and its copy, and use i,i' to refer to the original player and its copy.

The consistency assumption is that for any instance $I=(N,\mathcal{F})$ with bids b and its copy $I'=(N',\mathcal{F}')$ with bids b' such that for each pair i,i' we have $\max\{b'_i,b'_{i'}\}=b_i$, the number of winning bidders from each pair i,i' under b and b' is the same and the winner (if any) from each pair i,i' under b' is the bidder with the higher bid.

PROPOSITION 7.2. Consider a consistent deterministic pay-your-bid mechanism for a class of instances \mathcal{I} . Then for the smallest γ such that the underlying algorithm is γ -permeable for all instances $I \in \mathcal{I}$, there exists a pure Nash equilibrium that extracts at most a $1/\gamma$ fraction of the optimal social welfare.

Examples for both results can again be found in the context of single-minded combinatorial auctions. First note that combinatorial auctions naturally satisfy the requirement on the instances, we just have to duplicate players and let them demand the same bundle of items. Furthermore, both the algorithm that maximizes declared welfare as well as the two canonical greedy algorithms satisfy the consistency requirement.

8. CONCLUSION

In this paper we have considered the problem of characterizing algorithms that when combined with natural payment rules yield mechanisms with a low Price of Anarchy. For monotone, single-parameter algorithms we have identified a natural parameter, *permeability*, which provides a tight bound on the Price of Anarchy provable via smoothness. We have thus reduced the problem of characterizing the best possible smooth mechanism for a setting to the problem of characterizing the best possible permeability achieved by any algorithm for this setting.

We have used this reduction to provide the first lower bounds on the Price of Anarchy provable via smoothness, establishing the optimality of deterministic greedy algorithms for single-minded combinatorial auctions and of a novel randomized mechanism for interval graphs. These lower bounds apply across all algorithms, and, in particular, do not rely on polynomial-time computability.

Interesting open problems include (a) the extension of our results to non-monotone, single-parameter algorithms and (b) to provide a similar characterization for multiparameter mechanisms.

Acknowledgements

We would like to thank Eva Tardos for her valuable feedback.

REFERENCES

- BABAIOFF, M., LUCIER, B., NISAN, N., AND LEME, R. P. 2014. On the efficiency of the walrasian mechanism. In *Proceedings of the 15th Conference on Economics and Computation*. 783–800.
- BHAWALKAR, K. AND ROUGHGARDEN, T. 2011. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Symposium on Discrete Algorithms*. 700–709.
- BIKHCHANDANI, J. S., CHATTERJI, S., LAVI, R., MUALEM, A., NISAN, N., AND SEN, A. 2006. Weak monotonicity characterizes deterministic dominant strategy implementation. *Econometrica* 74, 1109–1132.
- Christodoulou, G., Kovács, A., and Schapira, M. 2008. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloquium on Automata, Languages and Programming.* 820–832.
- DÜTTING, P., HENZINGER, M., AND STARNBERGER, M. 2013. Valuation compressions in VCG-based combinatorial auctions. In *Proceedings of the 9th Conference on Web, Internet and Network Economics*. 146–159.
- DÜTTING, P., KESSELHEIM, T., AND TARDOS, E. 2015. Algorithms as mechanisms: The price of anarchy of relax-and-round. In *Proceedings of the 16th Conference on Economics and Computation*. Forthcoming.
- EDELMAN, B., OSTROVSKY, M., AND SCHWARZ, M. 2007. Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review 97*, 1, 242–259.
- FELDMAN, M., Fu, H., Gravin, N., and Lucier, B. 2013. Simultaneous auctions are (almost) efficient. In *Proceedings of the 46th Symposium on Theory of Computing*. 201–210.

- HARTLINE, J. D., HOY, D., AND TAGGART, S. 2014. Price of anarchy for auction revenue. In *Proceedings of the 15th Conference on Economics and Computation*. 693–710.
- JEHIEL, P. AND MOLDOVANU, B. 2001. Efficient design with interdependent valuations. *Econometrica* 69, 12371259.
- JEHIEL, P., MOLDOVANU, B., AND STACCHETTI, E. 1996. How (not) to sell nuclear weapons. *American Economic Review* 86, 814829.
- KRISHNA, V. AND MAENNER, E. 2001. Convex potentials with an application to mechanism design. *Econometrica* 69, 1113–1119.
- KULKARNI, J. AND MIRROKNI, V. S. 2015. Robust price of anarchy bounds via LP and fenchel duality. In *Proceedings of the 26th Symposium on Discrete Algorithms*. 1030–1049.
- LEHMANN, D. J., O'CALLAGHAN, L., AND SHOHAM, Y. 2002. Truth revelation in approximately efficient combinatorial auctions. *Journal of the ACM* 49, 5, 577–602.
- LUCIER, B. AND BORODIN, A. 2010. Price of anarchy for greedy auctions. In *Proceedings of the 21st Symposium on Discrete Algorithms*. 537–553.
- MCAFEE, R. P. AND MCMILLAN, J. 1988. Multi-dimensional incentive compatibility and mechanism design. *Journal of Economic Theory* 46, 335–354.
- MILGROM, P. 2004. Putting Auction Theory to Work. Cambridge University Press.
- MYERSON, R. 1981. Optimal auction design. *Mathematics of Operations Research* 6, 58–73.
- NISAN, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V. V. 2007. Algorithmic Game Theory. Cambridge University Press.
- ROCHET, J. 1987. A necessary and sufficient condition for rationalizability in a quasilinear context. *Journal of Mathematical Economics 16*, 191200.
- ROUGHGARDEN, T. 2012a. Intrinsic robustness of the price of anarchy. *Communications of the ACM 55*, 7, 116–123.
- ROUGHGARDEN, T. 2012b. The price of anarchy in games of incomplete information. In *Proceedings of the 13th Conference on Economics and Computation*. 862–879.
- ROUGHGARDEN, T. 2014. Barriers to near-optimal equilibria. In *Proceedings of the* 55th Symposium on Foundations of Computer Science. 71–80.
- SAKS, M. AND YU, L. 2005. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the 6th Conference on Electronic Commerce*. 286–293.
- SYRGKANIS, V. AND TARDOS, É. 2013. Composable and efficient mechanisms. In *Proceedings of the 45th Symposium on Theory of Computing*. 211–220.
- VARIAN, H. R. 2007. Position auctions. International Journal of Industrial Organization 25, 6, 1163–1178.