Combinatorial Auctions with Restricted Complements

Ittai Abraham

Microsoft Research, Silicon Valley
ittaia@microsoft.com

Moshe Babaioff

Microsoft Research, Silicon Valley
moshe@microsoft.com

Shaddin Dughmi

Microsoft Research, Redmond
shaddin@microsoft.com

Tim Roughgarden

Stanford University
tim@theory.stanford.edu

Complements between goods — where one good takes on added value in the presence of another — have been a thorn in the side of algorithmic mechanism designers. On the one hand, complements are common in the standard motivating applications for combinatorial auctions, like spectrum license auctions. On the other, welfare maximization in the presence of complements is notoriously difficult, and this intractability has stymied theoretical progress in the area. For example, there are no known positive results for combinatorial auctions in which bidder valuations are multi-parameter and non-complement-free, other than the relatively weak results known for general valuations.

To make inroads on the problem of combinatorial auction design in the presence of complements, we propose a model for valuations with complements that is parameterized by the "size" of the complements. The model permits a succinct representation, a variety of computationally efficient queries, and non-trivial welfare-maximization algorithms and mechanisms. Specifically, a *hypergraph-r valuation* v for a good set M is represented by a hypergraph $H = (M, E)$, where every (hyper-)edge $e \in E$ contains at most r vertices and has a nonnegative weight w_e. Each good $j \in M$ also has a nonnegative weight w_j. The value $v(S)$ for a subset $S \subseteq M$ of goods is defined as the sum of the weights of the goods and edges entirely contained in S.

We design the following polynomial-time approximation algorithms and truthful mechanisms for welfare maximization with bidders with hypergraph valuations.

(1) For bidders whose valuations correspond to subgraphs of a known graph that is planar (or more generally, excludes a fixed minor), we give a truthful and $(1 + \epsilon)$-approximate mechanism.
(2) We give a polynomial-time, r-approximation algorithm for welfare maximization with hypergraph-r valuations. Our algorithm randomly rounds a compact linear programming relaxation of the problem.
(3) We design a different approximation algorithm and use it to give a polynomial-time, truthful-in-expectation mechanism that has an approximation factor of $O(\log^r m)$.

1. INTRODUCTION

In a *combinatorial auction*, a set M of m heterogeneous and indivisible goods are allocated to n bidders with private preferences. We model the preferences of a bidder i via a *valuation* v_i, a function from 2^M to $\mathbb{R}^+$ with $v_i(\emptyset) = 0$ and $v_i(S) \leq v_i(T)$ for any

$S \subseteq T \subseteq M$. The number $v_i(S)$ indicates bidder i's "maximum willingness to pay" for the set of goods $S \subseteq M$. Several real-world scenarios can be modeled as combinatorial auctions, most famously FCC auctions for spectrum; see the books [Cramton et al. 2006; Milgrom 2004] for many more details and examples. One fundamental optimization problem in combinatorial auctions is *welfare maximization*, which is the problem of allocating the goods to the bidders to maximize $\sum_{i=1}^{n} v_i(S_i)$, where S_i denotes the goods allocated to i.

The difficulty of welfare maximization, and of combinatorial auction design more generally, depends on the structure of bidders' valuations. A key issue is whether valuations exhibit *complements*, where one good (e.g., a left shoe) has additional value in conjunction with another good (a right shoe). Formally, a *complement-free* (or *subadditive*) valuation v satisfies $v(S)+v(T) \geq v(S \cup T)$ for every pair $S, T \subseteq M$ of bundles. Most previous work in theoretical computer science on welfare maximization for combinatorial auctions has focused on complement-free valuations and several of its interesting subclasses (see Section 1.2 for details). A simple valuation v with complements is a *single-minded* one, where there is a bundle $S \subseteq M$ of at least two goods and a number $c > 0$ such that $v(T)$ equals c if $T \supseteq S$ and 0 otherwise.

Complements are prevalent in the motivating applications for combinatorial auctions. For example, in the FCC spectrum action there are natural synergies between nearby licenses. However, welfare maximization in the presence of complements is notoriously difficult, and this intractability has stymied theoretical progress in the area. For example, there are no known positive results for combinatorial auctions in which bidder valuations are multi-parameter and non-complement-free, other than the relatively weak results known for general valuations. Even with mere single-minded valuations, there is no polynomial-time approximation algorithm for the welfare maximization problem with sub-polynomial approximation ratio (assuming $P \neq NP$) [Lehmann et al. 2002; Sandholm 2002].

The combined importance and intractability of valuations with general complements motivates models for valuations with complements that are restricted in some way. Ideally, such a model should be parameterized by the "size" of the complements, and should permit a succinct representation, a variety of computationally efficient queries, and non-trivial welfare-maximization algorithms. In this paper we study such a model.

We study welfare maximization algorithms and truthful combinatorial auctions for a natural (hyper)graphical valuation model that concisely expresses limited complements. Precisely, a *hypergraph-r valuation* v for a good set M can, by definition, be represented by a hypergraph $H = (M, E)$, where every edge $e \in E$ contains at most r vertices and has a nonnegative weight w_e. Each good $j \in M$ also has a nonnegative weight w_j.[1] The value $v(S)$ for a subset $S \subseteq M$ of goods is defined as the sum of the weights of the goods and edges entirely contained in S:

$$v(S) = \sum_{j \in S} w_j + \sum_{e \,:\, e \subseteq S} w_e. \tag{1}$$

The parameter r is called the *rank* of the valuation v. The rank of a set of valuations is the maximal rank of any valuation in the set.

For example, consider a *graph* (i.e., hypergraph-2) valuation. A bidder with such a valuation has a base value for each good, and also enjoys bonuses if it acquires certain pairs of complementary goods. Even in this special case, each bidder possesses a polynomial number of different private parameters. Single-minded bidder valuations are essentially a single-parameter special case of hypergraph valuations — they

---

[1] No interesting approximation results are possible if weights are allowed to be negative (assuming $P \neq NP$).

correspond to the hypergraph valuations with a single edge with non-zero weight. A hypergraph-r valuation can be described via $O(m^r)$ parameters, which is polynomial in the number of goods when r is a fixed constant.

Summarizing, complements are an important, unavoidable aspect of combinatorial auctions, and they have been underrepresented in algorithmic mechanism design. Hypergraph valuations are a natural combinatorial representation of valuations that expresses complements in a controllable way. Even the most basic questions about welfare maximization with bidders with hypergraph valuations were completely open before the present work.

(1) What is the best-possible polynomial-time approximation ratio for welfare maximization (assuming $P \neq NP$), as a function of the valuations' ranks?
(2) Can equally good approximation ratios be achieved by incentive-compatible auctions?
(3) Does the combinatorial structure of the valuations permit novel combinatorial auction designs, for example by applying graph-theoretic tools?

1.1. Our Results

We contribute to all three of the questions above. Our main results are following.

(1) We use tools from graph theory to give a truthful and $(1 + \epsilon)$-approximate mechanism for bidders with valuations that correspond to subgraphs of a known graph that is planar (or more generally, excludes a fixed minor). Such valuations plausibly model some interesting combinatorial auctions, such as those where the predominant synergies involve "neighboring" goods (e.g., spectrum licenses).
Our technical approach here is to use Baker-type graph decompositions to design mechanisms that reduce the given welfare maximization problem to one with valuations supported by a graph with bounded treewidth, which we solve optimally in polynomial time using dynamic programming and the VCG mechanism.
(2) We give a polynomial-time r-approximation algorithm for welfare maximization with general hypergraph-r valuations. This guarantee is nearly the best possible, assuming $P \neq NP$. Our algorithm uses a correlated randomized rounding of a compact linear programming relaxation of the problem.
(3) With rich multi-parameter valuations like hypergraph valuations, only highly restricted types of approximation algorithms lead to truthful mechanisms; in particular, the r-approximation algorithm above does not. We instead design another algorithm, which randomly rounds the exponentially-sized configuration linear programming relaxation, and prove that it is "maximal-in-distributional-range (MIDR)" and hence induces a truthful-in-expectation mechanism. Our mechanism runs in polynomial time and has an approximation factor of $O(\log^r m)$ to the social welfare, for general hypergraph-r valuations. For small r, this guarantee is much better than the $O(\sqrt{m})$ approximation factor achieved by combinatorial auctions for general valuations, the only previous guarantee known for multi-parameter non-complement-free valuations.

1.2. Related Work

Combinatorial auctions are a paradigmatic problem in Algorithmic Mechanism Design [Nisan and Ronen 2001], as they are representative of problems in which incentives and computational constraints clash. For background on truthful approximation mechanisms for combinatorial auctions we refer the reader to [Cramton et al. 2006; Blumrosen and Nisan 2007]. As the literature on the subject is very large we only survey papers that are closely related to ours.

The most basic form of a valuation that exhibits complements is that of single minded bidders, in which each bidder desires one specific bundle. Lehmann, O'Callaghan and Shoham [2002] have presented a truthful mechanism that is an $O(\sqrt{m})$-approximation when bidders are single minded. It has also been shown [Sandholm 2002; Lehmann et al. 2002] that this is essentially the best possible (unless $P = NP$), even without incentive constraints. Lavi and Swamy [2005] have presented a randomized mechanism that is truthful-in-expectation, and achieves $O(\sqrt{m})$-approximation for general valuations. While this approximation is essentially the best possible for general valuations, it is obviously not attractive when m is large. Many previous works have sought improved approximations for restricted valuations; see [Blumrosen and Nisan 2007] for a survey. These have focused on subadditive valuations, and subclasses of them like submodular and "XOS" valuations. As far as we know, this paper is the first to improve the $O(\sqrt{m})$-approximation results for valuations that exhibit any form of complements.

Conitzer, Sandholm and Santi [2005] have considered "k-wise Dependent Valuations" which are similar to hypergraph-k valuations (except they also allow the edges' weights to be negative) and have shown that the problem of finding the efficient allocation is NP-hard even for graphs ($k = 2$), and have also studied the problem of eliciting such valuations. Chevaleyre et al. [2008] consider the same class of valuations (which they call "k-additive") and prove similar computational hardness results, and also study negotiation protocols between the agents. Unlike these papers our focus is on approximating the welfare-maximizing allocation and designing truthful mechanisms with good approximations.

Each of our technical results includes some ingredients that were developed in earlier works. Our r-approximation algorithm for welfare maximization (Theorem 4.1) is based on a sequence of randomized rounding algorithms devised for the multiway cut problem [Calinescu et al. 2000], the metric labeling problem [Kleinberg and Tardos 2002], graph homomorphism problems [Langberg et al. 2006], and other decomposition problems [Krauthgamer and Roughgarden 2011]. Our $O(\log^r m)$-approximate truthful-in-expectation combinatorial auction uses a number of ideas. The configuration linear programming relaxation, and the fact that its tractability reduces to that of demand queries, is folklore (see e.g. [Blumrosen and Nisan 2007]). The modified version we use with "proxy bidders" is based on an approach proposed by Feige [2009], and recently used by Dobzinski, Fu, and Kleinberg [2010] to obtain truthful mechanisms. To round this linear program we apply the decomposition technique of Lavi and Swamy [2005], which builds on Carr and Vempala [2002]. To argue truthfulness (in expectation) we use the concept of a "maximal-in-distributional-range (MIDR)" algorithm, which was first articulated by Dobzinski and Dughmi [2009], generalizing the "maximal-in-range (MIR)" definition in Nisan and Ronen [2007]. Finally, for our truthful approximation scheme for valuations supported by an excluded-minor graph, uses the decomposition theorem of Devos et al. [2004], which in turn uses the structure theorem of Robertson and Seymour [2004].

2. MECHANISM DESIGN PRELIMINARIES

We study direct-revelation mechanisms for players with quasi-linear utilities. In this setting, a *mechanism* is a pair $(\mathbf{X}, \mathbf{P})$ of algorithms. The *allocation rule* $\mathbf{X}$ takes as input a reported hypergraph valuation b_i from each player, and outputs a feasible assignment of goods to the players. The *payment rule* $\mathbf{P}$ takes the same input and computes a payment to the mechanism from each player. If a player i with a valuation v_i is assigned the goods $X_i(b_1, \ldots, b_n) \subseteq M$ and payment $P_i(b_1, \ldots, b_n) \geq 0$, then i earns *utility* $v_i(X_i(b_1, \ldots, b_n)) - P_i(b_1, \ldots, b_n)$. Let $x = X(v_1, \ldots, v_2)$ be the allocation when

all players report truthfully. The *welfare* $v(x)$ of the allocation $x = (x_1, \ldots, x_n)$ with respect to the valuations $v = (v_1, \ldots, v_n)$ is defined to be $v(x) = \sum_{i=1}^{n} v_i(x_i)$.

A deterministic mechanism is *truthful* if, for every player i and fixed reports b_{-i} by the players other than i, player i maximizes its utility by reporting its true valuation v_i,.

The *VCG mechanism* is a truthful mechanism that achieves the maximum-possible welfare (but generally not in polynomial time [Nisan and Ronen 2001]). This mechanism computes the allocation that maximizes the welfare with respect to the reported valuations, and then charges suitable "externality payments" to achieve truthfulness. A *maximal-in-range (MIR)* allocation rule pre-commits to a subset of feasible allocations, before receiving players' reports, and maximizes the welfare with respect to the reported valuations over this subset. Every MIR allocation rule induces a deterministic truthful mechanism via VCG-type payments [Nisan and Ronen 2007]. MIR allocation rules are useful when the subset of allowable allocations can be made structured enough to permit polynomial-time optimization, yet large enough to always contain a near-optimal allocation. A *maximal-in–distributional-range (MIDR)* allocation rule is the following randomized analogue: pre-commit to a set of distributions over feasible allocations, prior to receiving players' reported valuations; choose the distribution that maximizes expected welfare with respect to the reported valuations; and finally, sample a single feasible allocation from the chosen distribution. Every MIDR allocation rule induces a truthful-in-expectation mechanism when supplemented with VCG-type payments [Dobzinski and Dughmi 2009].

3. TRUTHFUL APPROXIMATION MECHANISMS FOR EXCLUDED-MINOR GRAPHS

In this section we impose restricted structure on bidders' hypergraph valuations and rewarded with a near-optimal truthful mechanism. The key assumption is that bidders valuations have rank 2, with the edges drawn from a known graph that excludes a fixed minor, such as a planar graph. This assumption is not unreasonable in the motivating applications for combinatorial auctions. For example, Kagel, Lien, and Milgrom [2010] have modeled player valuations in the FCC spectrum action as subgraphs of a fixed planar graph that captures the synergies of winning nearby licenses (similar models were suggested in Brunner et al. [2010]). In Appendix A we also present a truthful mechanism with approximation factor equal to the chromatic number of the square of the line graph of the graph.

3.1. Standard Definitions

We begin by recalling some standard definitions. A graph X is a *minor* of a graph G if there exists a sequence of edge deletions and edge contraction that start with G and end with X. A graph G *excludes X as a minor* if X is not a minor of G.

Definition 3.1. A graph $G = (M, E)$ has a *tree decomposition* with width k if there exists a tree (X, T), where $X = \{X_1, \ldots, X_n\}$ is a set of *bags* (where each $X_i \subseteq M$), and T is a tree whose nodes are the bags X_i. Such that

(1) $\bigcup X_i = M$ (every node in M appears in at least one bag);
(2) $\max |X_i| \le k + 1$ (each bag contains at most $k + 1$ nodes);
(3) for every $(u, v) \in E$ there exists i such that $u, v \in X_i$ (for every edge there exists a bag that contains both nodes);
(4) if X_i, X_j and X_ℓ are nodes, and X_ℓ is on the path in T from X_i to X_j, then $X_i \cap X_j \subseteq X_\ell$ (for every node in M, the set of bags that contain it induces a subtree of T).

The *treewidth* of a graph G is the least integer k for which G has a tree decomposition with width k.

A graph valuation (M, E_v) of a player is a *subgraph* of a given graph $G = (M, E)$ of every edge $e \in E_v$ belongs to E, that is, $e \in E_v$ implies $e \in E$.

3.2. Mechanisms

We next show that if all players' valuations are subgraphs of a given graph with *constant* treewidth then welfare maximizing allocation, and thus the efficient and truthful VCG mechanism, can be computed in polynomial time.

THEOREM 3.2. *If the valuations of all players are subgraphs of a given graph G with constant treewidth k, then the welfare maximizing allocation can be computed in time $n^{O(k)}$.*

To prove the theorem, let (T, X) be a tree decomposition of G, such that all bags are of size $\leq k+1$. By Bodlaender [1996] such a tree decomposition can be computed in linear time. Let v_ℓ be the valuation of player $\ell \in N$. For any bag $X_i \subseteq M$ let $A' : X_i \to N$ be a allocation of items in X_i (we call A' a partial allocation). Given two (partial) allocations A_1, A_2 and a set $Y \subseteq M$ we say that A_1 *agrees* with A_2 on Y and write $A_1(Y) = A_2(Y)$ if $\forall y \in Y, A_1(y) = A_2(y)$. We prove a slightly stronger statement using a standard dynamic programming induction argument.

LEMMA 3.3. *For any bag X_i and any partial allocation $A' : X_i \to N$, one can find*

$$\operatorname*{argmax}_{A | A(X_i) = A'(X_i)} \left\{ \sum_{\ell \in N} v_i(A^{-1}(\ell)) \right\}$$

the maximum efficient allocation that agrees with A' on X_i in time $n^{O(k)}$.

PROOF. Consider the tree T and fix an arbitrary root bag X_r. For any bag X_j let $M[X_j] \subseteq M$ be the set of all nodes that are in bags that are in the subtree of T that contains X_j and all its descendant bags (with respect to the root X_r).

Consider any two immediate children bags Y_1, Y_2 of a bag X_i and let $M_j = M[Y_j] \setminus X_i$ for $j \in \{1, 2\}$. The main observation is that there does not exists a pair $(u, v) \in M \times M$ such that $u \in M_1$ and $v \in M_2$ (this includes the case that $u = v$). This is true because by the property of a tree decomposition, such an edge (or vertex if $u = v$) would imply that both u and v must belong to X_i, but we removed X_i from both M_1 and M_2.

This observation suggests that the optimal allocation for $M[X_i]$ has the property that its restriction to $M[Y_j]$ is also the optimal solution for subgraph induced by $M[Y_j]$ (for any $j \in \{1, 2\}$). This implies that we can use a standard dynamic programming algorithm. Given a bag X_i we can compute for each allocation $A' : X_i \to N$ (there are at most n^k such allocations), the optimal allocation $A : M[X_i] \to N$ that agrees with A' (so $A'(X_i) = A(X_i)$). This can be done by going over all immediate children bags of X_i ($Y_1, \ldots, Y_\ell$) and for each such bag Y_j going over all partial allocations $A'_j : Y_j \to N$ that agree with $A' : X_i \to N$ and choose the one whose optimal allocation $A_j : M[Y_j] \to N$ for subgraph induced $M[Y_j]$ has maximal welfare. By the observation above, this procedure will find the optimal allocation for $M[X_i]$ that agrees with any partial allocation $A' : X_i \to N$. Moreover this can be done using $n^{k+O(1)}$ space and time using a standard dynamic programming approach. □

Using the result on bounded treewidth graphs, we devise an approximately efficient, truthful mechanism for valuation that are subgraphs of graphs excluding a fixed minor. We begin by presenting the result for planar graphs on which the basic tool we use is Baker's decomposition:

Common Knowledge: a planar graph $G = (M, E)$, parameter $\epsilon > 0$.
Input: reported graph valuations (that are subgraphs of G) $v_1 = (M, E_1, w_1), \ldots, v_n = (M, E_n, w_n)$.

(1) Let $k = \lceil 2/\epsilon \rceil$. Choose an arbitrary root r and for any $0 \le i \le k$ let P_i be the set of all points whose distance from r is $(k+1)\ell + i$ for some integer ℓ.
(2) For every $0 \le i \le k$ let M_i be the graph induced by $M \setminus P_i$. Note that M_i is k outer-planar and hence has tree width $\le 3k$.
(3) For every $0 \le i \le k$, compute the efficient allocation for M_i using the valuations that are induced by these items. (can be done in polynomial time by Theorem 3.2). Let A_i be the resulting allocation, and let R_i be the welfare of A_i.
(4) Pick any $i^* \in \operatorname{argmax}_i \{R_i\}$, then use the allocation A_{i^*}.

Fig. 1.   A maximal-in-range $(1+\epsilon)$-approximation algorithm when all players' valuations are subgraphs of a given planar graph

---

THEOREM 3.4. *[Baker 1994] For any planar graph $G = (M, E)$ and parameter k there exists a partition of G into $k + 1$ parts $P_0, \ldots, P_k$ such that for any $0 \le i \le k$ the graph induced by $M \setminus P_i$ has treewidth $\le 3k$.*

As shown in Baker [1994] a simple way to generate the elements of the partition is to choose an arbitrary root r and for any $0 \le i \le k$ let P_i be the set of all points whose distance from r is $(k+1)\ell + i$ for some integer ℓ.

THEOREM 3.5.   *Fix any $\epsilon > 0$. Assume that we are given a planar graph G such that the valuations of all the players are subgraphs of G. The algorithm of Figure 1, when combined with VCG payments, yields a truthful, polynomial time, $(1+\epsilon)$-approximation mechanism for the welfare maximization problem.*

PROOF.   Let $P_0, \ldots, P_k$ be the parts. For each graph G_i induced by the vertices $M \setminus P_i$, let A_i be the allocation that is maximally efficient for G_i among all allocations that are restricted to items $M \setminus P_i$. By Theorem 3.2 A_i can be computed in polynomial time. Note that the allocation rule that computes A_i is maximal-in-range, and therefore the algorithm that chooses the best A_i is also maximal in range. Maximal in range algorithms can be combined with truth-telling VCG payments, computable in polynomial time, to yield a truthful mechanism.

Consider an optimal allocation A. For any item $u \in M$ let β_u be the contribution of u to the welfare induced by A, for any edge $e = (u, v) \in E$ let β_e be the contribution of edge e to the welfare induced by A. Let $\beta = \sum_{u \in M} \beta_u + \sum_{e \in E} \beta_e$ be the optimal welfare obtained by A. For any part P_i, let α_i be the welfare of A from the items in P_i and the edges that have at least one vertex in P_i (formally $\alpha_i = \sum_{u \in P_i} \beta_u + \sum_{e \in E | e \cap P_i \ne \emptyset} \beta_e$). Observe that $\sum_{0 \le i \le k} \alpha_i \le 2\beta$ (since each vertex is counted once and each edge is counted at most twice). A mechanism that optimizes on the graph induced by $M \setminus P_i$ obtains welfare of at least $\beta - \alpha_i$. Thus the welfare of choosing the part P_i^* with the highest welfare is at least $\frac{1}{k+1} \sum_{0 \le i \le k} \beta - \alpha_i \ge \beta - \frac{2}{k}\beta \ge \beta(1-\epsilon)$.   □

We can extend Theorem 3.5 to a much larger family of graphs excluding any fixed minor using the following powerful decomposition theorem.

THEOREM 3.6. *[DeVos et al. 2004] For any graph X and parameter k there exists a constant $c = c(X, k)$ such that the following holds. For any graph $G = (M, E)$ that excludes X as a minor there exists a partition into $k + 1$ parts $P_0, \ldots, P_k$ such that for any i the graph induced by $M \setminus P_i$ has treewidth c.*

Moreover, the proof in DeVos et al. [2004] (along with the required structure theorem) implies that the partition can be computed in polynomial time. Using this and the same approach as in Theorem 3.5 gives the following result.

PROPOSITION 3.7. *Fix any $\epsilon > 0$ and a graph X. Assume that we are given a graph G that excludes X as a minor such that the valuations of all the players are subgraphs of G. There exists a truthful mechanism that runs in polynomial time and guarantees a $(1 + \epsilon)$-approximation to the social welfare.*

4. APPROXIMATE WELFARE MAXIMIZATION ALGORITHM

This section gives an r-approximation algorithm for welfare maximization with general and private hypergraph-r valuations. We do not expect that this algorithm, or small modifications to it, can lead to a truthful auction (see Remark 4.3). Nonetheless, this result makes two important points. First, in conjunction with known hardness results (Theorem 4.2), it precisely pins down the approximability of the welfare maximization problem for this class of bidder valuations. Second, it demonstrates that welfare maximization with hypergraph-r valuations is as tractable as with the far less expressive class of single-minded valuations with bundle size at most r.

Recall that an instance of the welfare maximization problem is described by n hypergraphs $H_1 = (M, E_1, w_1), \ldots, H_n = (M, E_n, w_n)$ of rank r, which represent players' valuations, all with a common vertex set M, which represents the goods. Feasible solutions are assignments of the goods to the players, which correspond to an ordered partition $S_1, \ldots, S_n$ of M. The goal is to compute the assignment maximizing the welfare $\sum_{i=1}^{n} v_i(S_i)$, where the hypergraph H_i defines the value $v_i(S_i)$ as in equation (1). We next present the main result of this section.

THEOREM 4.1. *There is a polynomial-time, r-approximation randomized algorithm for welfare maximization with hypergraph-r valuations.*

Our result almost matches the lower bound for the algorithmic problem.

THEOREM 4.2 ([TREVISAN 2001]). *There is no $r/2^{O(\sqrt{\log r})}$ approximation algorithm[2] for welfare maximization with hypergraph-r valuations, unless $P = NP$.*

Theorem 4.2 follows easily from hardness results for finding independent sets in bounded-degree graphs. It even applies to the special case of single-minded valuations with desired bundles of size at most r, which correspond to hypergraph-r valuations with only one edge with non-zero weight.

We next describe our approximation algorithm. Algorithms that use only "local" information to make allocation decisions, like natural greedy algorithms, do not seem capable of achieving an approximation ratio that depends only on the rank r.[3] Our algorithm solves and randomly rounds the following linear programming relaxation of the problem.

---

[2]In particular, for any $\epsilon > 0$ there is no $r^{1-\epsilon}$ approximation algorithm.

[3]This contrasts with single-minded valuations with bundle size at most r, for which a simple greedy algorithm provides an r-approximation. For example, consider an instance with 2 players and m goods. The first player has value $\sqrt{m}$, say, for the first good, and no value for anything else. The second player's valuation is a star centered on the first good, where each vertex has weight 0 and each edge has unit weight. The optimal solution gives all goods to the second player and has welfare $m - 1$, while most natural greedy algorithms allocate the first good to the first player and achieve welfare only $\sqrt{m}$.

$$\max \quad \sum_{i=1}^{n} \left(\sum_{j \in M} w_{ij} x_{ij} + \sum_{e \in E_i} w_{ie} z_{ie} \right) \qquad (2)$$

$$\text{subject to:} \qquad \sum_{i=1}^{n} x_{ij} = 1 \qquad \text{for every good } j. \qquad (3)$$

$$z_{ie} \leq x_{ij} \qquad \text{for every player } i, \text{ edge } e \in E_i,$$
$$\text{and good } j \in e. \qquad (4)$$

$$x_{ij} \geq 0 \qquad \text{for every player } i \text{ and good } j \in M. \quad (5)$$

$$z_{ie} \geq 0 \qquad \text{for every player } i \text{ and edge } e \in E_i. \quad (6)$$

Every feasible assignment naturally induces a 0-1 feasible solution to (2)–(6) with equal objective function value, where $x_{ij} = 1$ if and only if player i is assigned good j, and $z_{ie} = 1$ if and only if player i is assigned every good in e. The size of the linear program (2)–(6) is polynomial in the input size, so it can be solved in polynomial time.

We round the optimal solution to the linear program (2)–(6) using the randomized algorithm shown in Figure 2. The algorithm uses only the x-variables to make assignments; the z-variables are used in the analysis.

PROOF. **(of Theorem 4.1)** Since by equality (3) every good is fully assigned in x^*, the algorithm runs in polynomial time (both in expectation and with high probability). We next prove the approximation. We claim that, for every player i and edge $e \in E_i$, player i is assigned every good in e by the algorithm in Figure 2 with probability at least $z_{ie}^*/|e|$. Similarly, each player i is assigned good j with probability at least x_{ij}^*. This claim, combined with linearity of expectation and the assumption that $|e| \leq r$ for all i and $e \in E_i$, implies the theorem.

We observe that if a player ℓ is chosen in an iteration of the randomized rounding algorithm and no good of an edge e has been assigned in a previous iteration, then this player will get at least one good of e with probability $\max_{j \in e} x_{\ell j}^*$, and every good of e with probability $\min_{j \in e} x_{\ell j}^*$.

Now fix a player i and an edge e. (The same argument works for a good j, viewed as an edge of size 1.) We only need to lower bound the probability that i receives all of the goods in e in the same iteration of the randomized rounding algorithm.

Consider the first iteration q that assigns at least one good of e to some player. Using Bayes' rule and the fact that a single player is chosen uniformly at random each iteration, the probability that i is the chosen player in iteration q is

$$\frac{\max_{j \in e} x_{ij}^*}{\sum_{\ell=1}^{n} \max_{j \in e} x_{\ell j}^*} \geq \frac{\max_{j \in e} x_{ij}^*}{\sum_{\ell=1}^{n} \sum_{j \in e} x_{\ell j}^*} = \frac{\max_{j \in e} x_{ij}^*}{\sum_{j \in e} \sum_{\ell=1}^{n} x_{\ell j}^*} = \frac{\max_{j \in e} x_{ij}^*}{|e|}, \qquad (7)$$

where the final equality follows from the linear program constraints (3).

The probability that player i receives every good in e in round q, conditioned on i being chosen in iteration q (and hence receiving at least one good of e at this round), is

$$\frac{\min_{j \in e} x_{ij}^*}{\max_{j \in e} x_{ij}^*} \geq \frac{z_{ie}^*}{\max_{j \in e} x_{ij}^*}, \qquad (8)$$

where the inequality follows from the linear program constraints (4). Combining inequalities (7) and (8) shows that, independent of the value of q, player i is assigned every good in e with probability at least $z_{ie}^*/|e| \geq z_{ie}^*/r$. This completes the proof. $\square$

Input: an optimal solution (x^*, z^*) of the linear program (2)–(6).

(1) While there is at least one unassigned good:
    (a) Choose a player i uniformly at random.
    (b) Choose a threshold $t \in [0, 1]$ uniformly at random.
    (c) Assign to i every unassigned good j with $x_{ij}^* \geq t$.

Fig. 2. The randomized rounding algorithm for the linear program (2)–(6).

Remark 4.3. Turning the algorithm in Theorem 4.1 into a truthful mechanism with a similar approximation ratio seems difficult. Essentially, the only general technique for designing (randomized) truthful mechanisms for multi-parameter valuations like hypergraph valuations is via "maximal-in-distributional-range (MIDR)" approximation algorithms, which we define in the next section. The algorithm in Theorem 4.1 is not MIDR. There are several reasons for this; perhaps the most fundamental one is that hyperedges of different sizes lose different approximation factors in (7). The decomposition technique of Lavi and Swamy [2005] cannot be used to transform it into an MIDR algorithm, as our algorithm does not provide an approximation guarantee for hypergraphs with negative edge weights. (Negative edge weights cannot be ignored or rounded up to zero without changing the problem.) The convex rounding technique of Dughmi, Roughgarden, and Yan [2011] seems to yield interesting results only for subclasses of complement-free valuations, and not for valuations with complements as studied here.

5. A TRUTHFUL APPROXIMATION MECHANISM

In this section we present our truthful, $O(\log^r m)$-approximation mechanism for rank r valuations. Section 5.1 proves that demand oracles can be implemented in polynomial time for hypergraph valuations. Section 5.2 describes and analyzes our truthful approximate combinatorial auction for hypergraph valuations.

5.1. Demand Oracles and Hypergraph Valuations

The first step of our algorithm, described in the next section, solves an exponential-size linear programming relaxation of the welfare maximization problem. The complexity of solving this linear program reduces, by dualizing and the ellipsoid method, to that of a *demand oracle* for a valuation v: given a price p_j for each good $j \in M$, compute a bundle $S \in \arg\max\{v(S) - \sum_{j \in S} p_j\}$. Such queries can be answered in polynomial time for hypergraph valuations.

PROPOSITION 5.1. *Given a valuation v represented as a hypergraph $H = (M, E, w)$ and prices p for the goods of M, a bundle $S \subseteq M$ that maximizes $v(S) - \sum_{j \in S} p_j$ can be computed in polynomial time.*

PROOF. (Sketch.) First, a *value query* — given a subset $S \subseteq M$ of goods, return the value $v(S)$ — can be computed in time polynomial in the size of H by brute force, using equation (1).

Next, recall that a function v is supermodular if and only if, for every $S \subseteq M$ and $j, k \in M \setminus S$, $v(S \cup \{j, k\}) - v(S \cup \{j\}) \geq v(S \cup \{k\}) - v(S)$. Since edge weights are nonnegative, the definition in equation (1) of a hypergraph valuation v easily implies that $v(S)$, and hence also $v(S) - \sum_{j \in S} p_j$, is a supermodular function on M.

Finally, recall that every supermodular function can be maximized using a polynomial number of value queries (e.g. [Fleischer 2000]). □

5.2. Description and Analysis of the Mechanism

The main result of this section is the following theorem

THEOREM 5.2. *Fix any constant r. There is a polynomial time, truthful-in-expectation mechanism that guarantees $O(\log^r m)$-approximation to the social welfare when all players have hypergraph-r valuations.*

We now describe our new allocation rule. Like the welfare maximization algorithm in Section 4, it randomly rounds a linear programming relaxation. To obtain both truthfulness-in-expectation and a good approximation, we make two non-trivial changes. First, our linear program fractionally allocates each good at most B times, where $B = \Omega(\log m)$. On one hand, this duplication of the items results in a linear program with a constant integrality gap. On the other hand, rounding the linear program involves more "conflict resolution". We define a natural rounding procedure that degrades the objective value by a factor of at most $O(B^r)$. The objective function of our linear program will involve modified valuation profiles, first proposed by Feige [2009] and termed *proxy valuations* by Dobzinski, Fu, and Kleinberg [2010]. The proxy valuations will take into account the rounding procedure, and guarantee that the allocation rule that first solves the linear program and then rounds the fractional solution is MIDR.

Our high-level algorithm is described in Figure 3. In more detail, the first step is to scale the edge weights of the reported valuations $v_1, \ldots, v_n$ to obtain the corresponding proxy valuations $v'_1, \ldots, v'_n$, as follows,

$$v'_i(S) = \sum_{v \in S} \frac{w_{iv}}{B} + \sum_{e \,:\, e \subseteq S} \frac{w_{ie}}{B^{|e|}}. \tag{9}$$

We call v'_i the *proxy valuation* corresponding to v_i. Note that the proxy valuation of any set is exactly the expected value of the set when each item is assigned to the bidder independently with probability $1/B$. The next step formulates the following linear program

$$\max \ f(\mathbf{y}) = \sum_{i, S \neq \emptyset} v'_i(S) y_{i,S} \tag{10}$$

$$\text{subject to:} \quad \sum_{S \neq \emptyset} y_{i,S} \leq 1 \qquad \text{for every player } i. \tag{11}$$

$$\sum_i \sum_{S | j \in S} y_{i,S} \leq B \quad \text{for every good } j \in M. \tag{12}$$

$$y_{i,S} \geq 0 \qquad \text{for every player } i \text{ and bundle } S \subseteq M. \tag{13}$$

While the linear program has exponential size, it can be solved in polynomial time (by dualizing and using the ellipsoid algorithm) given access to a demand oracle [Blumrosen and Nisan 2005]. By Proposition 5.1, such a demand oracle can be computed in polynomial-time for hypergraph-r valuations. The next step is to use the Lavi-Swamy [Lavi and Swamy 2005] approach to compute, in polynomial time, a distribution over integral solutions with polynomial size support with the following property. The expectation of this distribution over integer solutions is equal to the optimal fractional solution, divided by a universal constant α. Any choice of α greater than the integrality gap of the linear program suffices; The integrality gap is $m^{\frac{1}{B+1}}$, which is bounded by a universal constant for any choice of $B = \Omega(\log m)$ (c.f. [Lavi and Swamy 2005]). In the fourth step we choose one of the integral solutions according to the prescribed distribution. Finally, in the last stage we resolve conflicts in the integral solution in such a way that if a player received a copy of the good in the integral solution then he is allocated that good with probability $1/B$, independently over all goods.

Parameters: $B = \Omega(\log m)$. Universal constant α exceeding integrality gap of linear program (10)–(13).
Input: reported hypergraph-r valuations $v_1 = (M, E_1, w_1), \ldots, v_n = (M, E_n, w_n)$.

(1) Form proxy valuations v_i' from v_i for every player i. Let $v' = (v_1', v_2', \ldots, v_n')$.
(2) Solve the linear program (10)–(13) with respect to v', obtaining optimal solution $\mathbf{y}^* = \{y_{iS}^*\}_{1 \leq i \leq n, S \subseteq M}$.
(3) Compute a decomposition $\frac{\mathbf{y}^*}{\alpha} = \sum_{\ell \in \mathbf{L}} \lambda_l \mathbf{z}_\ell$ into a convex combination of a polynomial number of feasible integral solutions.
(4) Sample a feasible integral solution $\mathbf{z} = \{z_{iS}^*\}_{1 \leq i \leq n, S \subseteq M}$ according to the distribution $\{\lambda_\ell\}_{\ell \in \mathbf{L}}$.
(5) Independently for each good $j = 1, 2, \ldots, m$: let $I_j = \{i \mid z_{i,S} = 1, j \in S\}$, then with probability $\frac{|I_j|}{B}$ allocate j to a uniformly random player in I_j, otherwise do not allocate j. Note that $|I_j| = \sum_i \sum_{S|j \in S} z_{i,S} \leq B$ for every good j, thus $\frac{|I_j|}{B} \leq 1$.

Fig. 3. The MIDR randomized rounding allocation rule for the linear program (10)–(13).

From the above description and Proposition 5.1 is it clear that the mechanism runs in polynomial time. Next we prove the approximation guarantee and show that the mechanism is MIDR.

PROPOSITION 5.3. *The mechanism is MIDR and the approximation guarantee is at most αB^r*

PROOF. First, we prove that the mechanism is MIDR. Let $M(v)$ denote the expected social welfare achieved by the mechanism on receiving reports v. Let $\mathbf{x}(\mathbf{z})$ be the integral allocation of the mechanism (at the end of stage (5)) given a feasible integral solution $\mathbf{z}$ from stage 4. Then from linearity of expectation and the fact that per player, the probability of getting allocated each good in step (5) is independent we have

$$\mathbf{E}[v(\mathbf{x}(\mathbf{z}))] = \sum_{i, S \neq \emptyset} v_i'(S) z_{i,S} = f(\mathbf{z}). \tag{14}$$

for any $\mathbf{z}$. By definition of the decomposition of stage 3, and using linearity of $f(.)$, we have

$$\mathbf{E}_{\{\lambda_\ell\}_{\ell \in \mathbf{L}}}[f(\mathbf{z}_\ell)] = f(\mathbf{y}^*)/\alpha. \tag{15}$$

Thus, from equalities (14),(15) we conclude that

$$M(v) = \mathbf{E}_{\{\lambda_\ell\}_{\ell \in \mathbf{L}}}[\mathbf{E}[v(\mathbf{x}(\mathbf{z}_l))]] = f(\mathbf{y}^*)/\alpha \tag{16}$$

Since the linear program maximizes f over its feasible domain, equation (16) implies that the mechanism optimizes M over the set of distributions in its range. Thus, the mechanism is MIDR.

We now prove the approximation guarantee. Let x^* denote the vector encoding the welfare-maximizing allocation for valuation profile v, let $v(x^*)$ denote the optimum welfare, and $v'(x^*)$ denote the corresponding "proxy" welfare. Observe that $f(y^*) \geq f(x^*)$ by optimality of y^* for the linear program. Since x^* is a feasible integral allocation, $f(\mathbf{x}^*) = v'(\mathbf{x}^*)$. By equation (9) we have that

$$v'(\mathbf{x}) \geq \frac{v(\mathbf{x})}{B^r} \tag{17}$$

for any integral allocation x, in particular, for $\mathbf{x}^*$. Combining with (16),(17) it follows that

$$M(v) = \frac{f(\mathbf{y}^*)}{\alpha} \geq \frac{v'(\mathbf{x}^*)}{\alpha} \geq \frac{v(\mathbf{x}^*)}{\alpha B^r} \tag{18}$$

□

REFERENCES

BAKER, B. S. 1994. Approximation algorithms for np-complete problems on planar graphs. *J. ACM 41,* 1, 153–180.

BLUMROSEN, L. AND NISAN, N. 2005. On the computational power of iterative auctions. In *ACM EC*. 29–43.

BLUMROSEN, L. AND NISAN, N. 2007. *In Algorithmic Game Theory (Noam Nisan, Tim Roughgarden, Eva Tardos and Vijay Vazirani (Editors)). Combinatorial Auctions (a survey)*. Cambridge University Press.

BODLAENDER, H. L. 1996. A linear-time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.*, 1305–1317.

BRUNNER, C., GOEREE, J. K., HOLT, C. A., AND LEDYARD, J. O. 2010. An experimental test of flexible combinatorial spectrum auction formats. *American Economic Journal: Microeconomics 2,* 1, 39–57.

CALINESCU, G., KARLOFF, H. J., AND RABANI, Y. 2000. An improved approximation algorithm for multiway cut. *J. Comput. Syst. Sci. 60,* 3, 564–574.

CARR, R. AND VEMPALA, S. 2002. Randomized metarounding. *Random Struct. Algorithms 20*, 343–352.

CHEVALEYRE, Y., ENDRISS, U., ESTIVIE, S., AND MAUDET, N. 2008. Multiagent resource allocation in k-additive domains: Preference representation and complexity. *Annals of Operations Research 163,* 1, 49–62.

CONITZER, V., SANDHOLM, T., AND SANTI, P. 2005. Combinatorial auctions with k-wise dependent valuations. In *AAAI*. 248–254.

CRAMTON, P., SHOHAM, Y., AND STEINBERG, R. 2006. *Combinatorial Auctions*. MIT Press.

DEVOS, M., DING, G., OPOROWSKI, B., SANDERS, D. P., REED, B., SEYMOUR, P., AND VERTIGAN, D. 2004. Excluding any graph as a minor allows a low tree-width 2-coloring. *Journal of Combinatorial Theory, Series B 91,* 1, 25 – 41.

DOBZINSKI, S. AND DUGHMI, S. 2009. On the power of randomization in algorithmic mechanism design. In *FOCS*. 505–514.

DOBZINSKI, S., FU, H., AND KLEINBERG, R. 2010. Truthfulness via proxies. *CoRR abs/1011.3232*.

DUGHMI, S., ROUGHGARDEN, T., AND YAN, Q. 2011. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *STOC*. 149–158.

FEIGE, U. 2009. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput. 39*, 122–142.

FLEISCHER, L. 2000. Recent progress in submodular function minimization. Mathematical Programming Society Newsletter, OPTIMA64.

KAGEL, J., LIEN, Y., AND MILGROM, P. 2010. Ascending prices and package bidding: A theoretical and experimental analysis. *American Economic Journal: Microeconomics 2,* 3, 160–85.

KLEINBERG, J. M. AND TARDOS, É. 2002. Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields. *J. ACM 49,* 5, 616–639.

KRAUTHGAMER, R. AND ROUGHGARDEN, T. 2011. Metric clustering via consistent labeling. *Theory of Computing 7,* 1, 49–74.

LANGBERG, M., RABANI, Y., AND SWAMY, C. 2006. Approximation algorithms for graph homomorphism problems. In *APPROX-RANDOM*. 176–187.

LAVI, R. AND SWAMY, C. 2005. Truthful and near-optimal mechanism design via linear programming. In *FOCS*. 595–604.

LEHMANN, D. J., O'CALLAGHAN, L., AND SHOHAM, Y. 2002. Truth revelation in approximately efficient combinatorial auctions. *J. ACM 49,* 5, 577–602.

MILGROM, P. 2004. *Putting Auction Theory to Work*. Cambridge University Press.

NISAN, N. AND RONEN, A. 2001. Algorithmic Mechanism Design. *Games and Economic Behavior 35,* 1-2, 166–196.

NISAN, N. AND RONEN, A. 2007. Computationally feasible VCG mechanisms. *J. Artif. Intell. Res. (JAIR) 29*, 19–47.

ROBERTSON, N. AND SEYMOUR, P. 2004. Graph minors. xx. wagner's conjecture. *Journal of Combinatorial Theory, Series B 92,* 2, 325 – 357. Special Issue Dedicated to Professor W.T. Tutte.

SANDHOLM, T. 2002. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence 135,* 1-2, 1 – 54.

TREVISAN, L. 2001. Non-approximability results for optimization problems on bounded degree instances. In *STOC*. 453–461.

A. WHEN THE SQUARE OF THE LINE GRAPH HAS A SMALL CHROMATIC NUMBER

In this section we present a truthful mechanism with approximation that equals to the chromatic number of the square of the line graph of the common graph G (when all valuations are subgraphs of G). For example, if G has bounded degree d then the square of its line graph has chromatic number $O(d^2)$.

We begin with some standard definitions. Given a graph $G = (V, E)$ let $L(G) = (E, \{(u, v), (w, x) \in E \mid |\{u, v\} \cap \{w, x\} \neq \emptyset\})$ be the *line graph* of G whose vertices are the edges of G and nodes in $L(G)$ share an edge in $L(G)$ if the two corresponding edges in G intersect at some vertex of G. Given a graph $G = (V, E)$ let $G^2 = (V, \{(u, v) \mid \exists w, (u, w), (w, v) \in E \lor (u, v) \in E\})$ be the *square graph* whose edges are all length one and two paths in G.

We are now ready to present our result.

PROPOSITION A.1. *Let G by any graph such that the valuations of all players are subgraphs of G. Then there exists a truthful mechanism that runs in polynomial time and is a $\chi(L(G)^2)$-approximation.*

PROOF. One can color edges of the square graph with $\chi(L(G)^2)$ colors. So any two edges with the same color are at distance at least 2. We can iterate over all colors and for each color allocate only the items that are induced by the edges of that color. Since edges are at distance at least 2, we can compute the optimal allocation for each such edge independently. □

One simple corollary of the proposition is that if the common graph G has bounded degree d then we can design a truthful mechanism with $O(d^2)$-approximation.