Combinatorial Auctions with Tractable Winner Determination

GEORG GOTTLOB, University of Oxford, UK
and
GIANLUIGI GRECO, Università della Calabria, Italy

Some recent advances in the identification of tractable classes of combinatorial auctions are discussed. In particular, the work of [Gottlob and Greco 2007] is illustrated, where a research question raised in [Conitzer et al. 2004] is solved by showing that the class of *structured item graphs* is not efficiently recognizable (i.e., deciding the membership of instances is NP-hard), and where this difficulty is overcome trough a different approach based on the notion of *hypertree decomposition*.

Categories and Subject Descriptors: F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems

General Terms: Algorithms, Economics, Theory

1. STRUCTURAL APPROACHES FOR WINNER DETERMINATION

Combinatorial auctions are well-known mechanisms for resource and task allocation where bidders are allowed to simultaneously bid on combinations of items. This is desirable when a bidder's valuation of a bundle of items is not equal to the sum of her valuations of the individual items. Determining an allocation of the items among the bidders that maximizes the sum of the accepted bid prices is a crucial problem in combinatorial auctions (see, e.g., [Lehmann et al. 2006]). This problem, called *winner determination*, is known to be NP-hard [Rothkopf et al. 1998], and is not approximable in polynomial time unless NP = ZPP [Sandholm 2002]. Intensive efforts have been made to identify restricted classes of combinatorial auctions for which winner determination is tractable. In particular, such classes were obtained by identifying structural restrictions of the interaction among bidders that are likely to occur in practice. This avenue of research was first investigated in a systematic way by [Sandholm and Suri 2003] and by [Conitzer et al. 2004], where the notion of *item graph* was used to represent bidder interaction in graphical terms, and where additional "structural" restrictions imposed on item graphs are investigated to guarantee the polynomial time solvability of the winner determination problem.

In a nutshell, an item graph is a graph whose nodes are in one-to-one correspon-

Authors' addresses: Gianluigi Greco, Dipartimento di Matematica - Università della Calabria, Via Pietro Bucci 30B, 87030 Rende (CS), Italy, email: ggreco@mat.unical.it. Georg Gottlob, Oxford University Computing Laboratory, OX1 3QD Oxford, UK, email: georg.gottlob@comlab.ox.ac.uk.

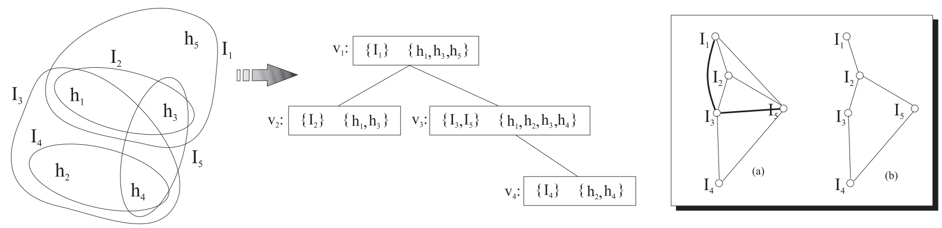

Fig. 1. Example Dual Hypergraph and Hypertree Decomposition; (a), (b) two item graphs.

dence with the items of an auction, and edges are such that for any bid, the items occurring in it induce a connected subgraph.

As an example, the graphs (a) and (b) in the rightmost part of Figure 1 are two different item graphs associated with the combinatorial auction over items $\{\mathrm{I}_1, ..., \mathrm{I}_5\}$ and where bids are on the following bundle of items: $h_1 : \{\mathrm{I}_1\}$, $h_2 : \{\mathrm{I}_1, \mathrm{I}_2, \mathrm{I}_3\}$, $h_3 : \{\mathrm{I}_1, \mathrm{I}_2, \mathrm{I}_5\}$, $h_4 : \{\mathrm{I}_3, \mathrm{I}_4\}$, and $h_5 : \{\mathrm{I}_3, \mathrm{I}_4, \mathrm{I}_5\}$.

An important parameter of a graph is its *treewidth* [Robertson and Seymour 1986], which is a measure of its degree of cyclicity. Restrictions on the treewidth of item graphs of combinatorial auctions were studied by [Conitzer et al. 2004], where item graphs whose treewidth is bounded by a constant are called *structured item graphs*. [Conitzer et al. 2004] have shown that the winner determination problem is tractable on classes of combinatorial auctions having structured item graphs. In fact, bidding interaction of an auction having an item graph with treewidth k can be represented as a tree whose vertices are clusters of at most $k+1$ items, the so-called *tree decomposition* of the item graph [Robertson and Seymour 1986]. This tree decomposition is exploited by [Conitzer et al. 2004] for deriving an efficient algorithm for winner determination.

Clearly enough, this tractability result is useful only if a structured item graph is given or can efficiently be determined. However, exponentially many item graphs might be associated with a given auction, in correspondence with the possible ways of preserving the connectivity condition—for instance, graph (b) in Figure 1 is obtained from graph (a) by deleting the arcs depicted in bold, but different simplifications are possible. While deciding whether a given (item) graph has treewidth bounded by a fixed natural number k is known to be feasible in linear time—in fact, in time $O(2^{ck^2} \times n)$ for graphs on n vertices [Bodlaender 1996]—, it was left as an open problem in [Conitzer et al. 2004] whether it is tractable to check if for a combinatorial auction, an item graph of treewidth bounded by a fixed natural number k exists and whether, if so, this item graph can be constructed in polynomial time.

2. HYPERTREE DECOMPOSITIONS FOR COMBINATORIAL AUCTIONS

The above research question has recently been answered by [Gottlob and Greco 2007] who proved that deciding whether an item graph of treewidth ≤ 3 exists is NP-hard. This is bad news.

Motivated by this result, [Gottlob and Greco 2007] subsequently investigated other structural restrictions that might be used to single out classes of tractable combinatorial auctions which are, moreover, efficiently recognizable. This time, the investigation leads to very good news. A novel class of tractable auctions has been singled out based on the following two key ideas:

(1) Bidder interaction is represented trough the *dual hypergraph*, whose nodes are the various bids, and that, for each item I in the auction, has a hyperedge consisting of the set of bids that contain item I. As an example, the hypergraph at the left of Figure 1 encodes the combinatorial auction discussed before—e.g., there is a hyperedge I_1 over the bundles h_1, h_3, and h_5.

(2) The structural intricacy of a dual hypergraph is measured using the notion of *hypertree width*, which is based on the concept of *hypertree decomposition* [Gottlob et al. 2002]. The latter is a method for appropriately transforming any hypergraph into an equivalent acyclic one by organizing its edges and nodes into a polynomial number of clusters, and by suitably arranging the clusters as a tree (see Figure 1). The more "cyclic" a hypergraph is, the bigger is the maximum cluster size. The minimum size required over all the possible decompositions is the hypertree width of the hypergraph. For a comparison of hypertree width with other measures for hypergraph cyclicity, see [Gottlob et al. 2000].

The hypertree approach has two important advantages. First, differently from structured item graphs, for any fixed constant k, deciding whether a hypergraph has hypertree width k and computing a hypertree decomposition of width k, if so, are known to be feasible in polynomial time. Second, for fixed k, those instances of the winner determination problem whose dual hypergraphs have hypertree width bounded by k are shown in [Gottlob and Greco 2007] to be efficiently solvable by means of a bottom-up procedure that incrementally constructs a solution by starting from the leaves of a hypertree decomposition: At each decomposition vertex v, all the possible solutions for the problem restricted over the bids occurring in the subtree rooted at v are computed (in an implicit and succinct way), based on the results available at v's children.

Thus, hypertree decompositions on dual hypergraphs represent a viable way for isolating efficiently recognizable as well as efficiently solvable instances.

There is further good news in [Gottlob and Greco 2007]. Indeed, [Gottlob and Greco 2007] also embark on a comparison with structured items graphs and show that nothing is lost in terms of generality when considering the hypertree decomposition of dual hypergraphs. To the contrary, it has been shown that strictly larger classes of instances are efficiently solvable according to the new approach than according to the structured item graphs approach.

In fact, structured item graphs turn out to be in one-to-one correspondence with special kinds of hypertree decompositions of the dual hypergraphs satisfying some additional requirements (which are conceptually close to the notion of *query decomposition* [Chekuri and Rajaraman 1997]). And, as a matter of fact, these requirements represent an obstacle for the efficient recognizability of the corresponding class of auctions, while they are actually not needed for ensuring the tractability of the winner determination problem. In fact, we currently do not know the precise relationship between hypertree width of dual hypergraphs and treewidth of associated item graphs. In particular, we do not know whether there is some scenario with "small" dual hypertree width but "large" treewidth for each associated item graph. Indeed, establishing this relationship requires a deep understanding of the approximability of (the related notion of) query decomposition, which is currently missing in the literature.

We conclude by noticing that assessing whether hypertree decompositions can be used to deal with generalizations of the winner determination problem (e.g, several copies of the same item are available and the auctioneer is satisfied when at least a given number of copies is actually sold) has not been considered by [Gottlob and Greco 2007]. This is an interesting avenue of further research.

REFERENCES

Bodlaender, H.L. 1996. A Linear-Time Algorithm for Finding Tree-Decompositions of Small Treewidth. *SIAM Journal on Computing 25(6)*, 1305–1371.

Chekuri, Ch. and Rajaraman, A. 1997. Conjunctive query containment revisited. In *Proc. of ICDT'07*. 56-70.

Conitzer, V., Derryberry, J., and Sandholm, T. 2004. Combinatorial auctions with structured item graphs. In *Proc. of AAAI'04*. 212–218.

Gottlob, G. and Greco, G. 2007. On The Complexity of Combinatorial Auctions: Structured Item Graphs and Hypertree Decompositions. In *Proc. of EC'07*. 152–161.

Gottlob, G., Leone, N., and Scarcello,F. 2000. A Comparison of Structural CSP Decomposition Methods. *Artificial Intelligence* 124:2, 243–282.

Gottlob, G., Leone, N., and Scarcello, F. 2002. Hypertree decompositions and tractable queries. *Journal of Computer and System Sciences 63(3)*, 579–627.

Lehmann, D., Müller, R., and Sandholm, T. 2006. The Winner Determination Problem. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*. MIT Press.

Robertson, N. and Seymour, P. 1986. Graph minors ii. algorithmic aspects of tree width. *Journal of Algorithms 7*, 309–322.

Rothkopf, M. H., Pekec, A., and Harstad, R. M. 1998. Computationally manageable combinatorial auctions. *Management Science 44*, 1131–1147.

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

Sandholm, T. and Suri, S. 2003. Bob: Improved winner determination in combinatorial auctions and generalizations. *Artificial Intelligence 7*, 33–58.