DIMACS Technical Report 95-09
April 1995

Computationally Manageable Combinatorial Auctions[1]

by

Michael H. Rothkopf[2]
School of Business
and
Rutgers Center for Operations Research
Rutgers University

Aleksandar Pekeč[3]
Department of Mathematics
Rutgers University

Ronald M. Harstad
Faculty of Management
and
Rutgers Center for Operations Research
Rutgers University

[1]The authors gratefully acknowledge the National Science Foundation for its support under grant number SBR 93-09333 to Rutgers University
[2]Permanent Member
[3]Affiliated Graduate Student Member, DIMACS Graduate Assistant

DIMACS is a cooperative project of Rutgers University, Princeton University, AT&T Bell Laboratories and Bellcore.

DIMACS is an NSF Science and Technology Center, funded under contract STC–91–19999; and also receives support from the New Jersey Commission on Science and Technology.

ABSTRACT

There is interest in designing simultaneous auctions for situations in which the value of assets to a bidder depends upon which other assets he or she wins. In such cases, bidders may well wish to submit bids for combinations of assets. When this is allowed, the problem of determining the revenue maximizing set of nonconflicting bids can be a difficult one. We analyze this problem, identifying several different structures of combinatorial bids for which computational tractability is constructively demonstrated and some structures for which computational tractability

1 Introduction

Some auctions sell many assets simultaneously. Often these assets, like U.S. treasury bills, are interchangeable. However, sometimes the assets and the bids for them are distinct. This happens frequently, as in the U.S. Department of the Interior's simultaneous sales of off-shore oil leases, in some private farm land auctions, and in the Federal Communications Commission's recent multi-billion dollar sales of the rights to use radio spectrum. It also happened in the post-World-War-II divestiture of synthetic rubber plants by the U.S. government. In such situations, it may well be that the value of an asset to a bidder depends strongly upon which other assets he wins. In off-shore oil-lease bidding, this dependency often takes the form of diseconomies of scale. (The diseconomies of scale are due to risks associated with the extreme variability of returns and the large amounts of money involved. See [18] for a discussion of bidding in simultaneous auctions with a constraint on exposure, i.e. the total of all bids. Such constraints are imposed on off-shore oil lease bidding teams by their managements to control risks.) However, in each of the other examples there are clearly situations in which the value of an asset is increased if another asset or group of assets is won. For example, in the radio spectrum auctions, a licence for the Philadelphia region may be much more valuable to a company if that company also has licences for the New York and/or the Washington/Baltimore regions.

Because of the possibility of such synergy or superadditivity in values, the designers of simultaneous sales have reason to consider allowing bids not just for individual assets, but for combinations of assets. Off-shore oil lease sales have not allowed such bids. Some farmland sales do ([20]) . As described by McCurdy [15] and Rothkopf [19], the divestiture process for synthetic rubber plant units was originally crafted by Congress without such bids, but when Shell Chemical Company under McCurdy submitted a single bid for three units that exceeded the sum of bids from others for those units, Congress voted to accept the Shell bid. The design of the FCC's spectrum auctions was subject to an unusually sophisticated debate. The FCC's original notice of proposed rule making ([5]), proposing an auction design, received many responses from telecommunications companies, a number of them with papers by leading economists appended. (Among those contributing were Paul Milgrom and Robert Wilson for Pacific Bell and Nevada Bell, Robert Weber for TDS, R. Preston McAfee for Airtouch, Robert Harris and Michael Katz for Nynex, Barry J. Nalebuff and Jeremy I. Bulow for Bell Atlantic, Peter Cramton for MCI, and R. Mark Isaac for CTIA. There were also sophisticated rounds of replies to the responses, three conferences to discuss issues in the auction design, and pilot experiments conducted at the California Institute of Technology. McMillan [14] describes many of the issues.

The potential advantage of allowing combinatorial bids in a simultaneous auction is that it allows bidders to express their synergistic values. When such synergies are present, this should result in both greater revenue for the bid-taker and in an economically more efficient allocation of assets to bidders. Bykowsky, Cull and Ledyard [3] illustrate situations in which permitting combinatorial bids is essential to an efficient outcome.

There are two potential disadvantages of allowing combinatorial bids. The first of these, the "threshold" problem (inaptly called the "free–rider" problem in Milgrom and Wilson [16]), can occur when diseconomies of scale predominate. For example, four bidders are competing for two assets: A and B. Suppose that each bidder's valuation is private information and that no bidder values any asset except as follows: A is worth \$100 to bidder 1; B is worth \$100 to bidder 2; and the combination AB is worth \$150 to bidder 3 and \$110 to bidder 4. In this situation, the most economic allocation is to give bidder 1 asset A and bidder 2 asset B. However, if the bids for individual assets A and B are each at \$1, neither of these bidders acting unilaterally can afford the \$110 necessary to match a bid of, say, \$111 for the combination AB by bidder 3. It is worth noting that the threshold problem has an opposing force, the "exposure" problem. This problem is that an unsuccessful attempt to acquire a collection of assets, when combinatorial bidding is not allowed, may lead to paying more for some individual assets than they are worth. Alternatively, a bidder unwilling to risk bidding above his individual valuations on individual assets may not be able to obtain a combination for which synergies make him the efficient recipient. This arises in the above example with a single change that makes AB worth \$250 to bidder 3. If synergies are important, the efficiency consequences of the exposure problem are likely to outweigh those of the threshold problem. Both problems can be magnified when many assets are sold and large combinations are sought.

The second potential difficulty with allowing combinatorial bids is the computational difficulty of finding the best, i.e. revenue maximizing, set of winning bids. This issue was raised in the FCC auction design debate by McAfee [12]. In the worst case, the bid taker offering n assets could receive offers on $2^n - 1$ different combinations of assets. Clearly, bid evaluation will present a computational problem when n is large. The FCC is in the process of selling over 2,500 licences. While not all of these are being sold simultaneously, an auction just completed sold 99 major licences for just over $7 billion. The FCC plans future simultaneous sales involving several hundred (less valuable) licences.

This paper considers the computational problem of evaluating combinatorial bids. It first gives some additional background including a more detailed description of the FCC auction and a discussion of alternative formats for simultaneous sale of many assets. Then, after discussing the problem of evaluating unrestricted combinatorial bidding, it develops various structures of permitted combinatorial bids for which the computational problems are provably manageable even for the simultaneous sale of many assets. In a concluding section, we argue that allowing restricted combinatorial bidding may provide many of the potential advantages of unrestricted combinatorial bids without imposing an impractical computational burden.

A complete analysis of the desirability of certain structures of allowable combinatorial bids would have to go further than this paper. Manageability of winner determination is, however, the logical first step. Consideration of the revenue impact of permitting some combined bids depends upon an analysis of how allowing combinatorial bidding affects bidders' behavior. A necessary first step toward modeling bidder behavior is determining how large a bid on a particular combination must be to avoid losing, given a current set of competing bids for overlapping combinations. We show below that the manageability of this problem of determining a minimal improving (i.e., revenue–increasing) combinatorial bid is essentially equivalent to the bid–taker's problem that is our focus.

2 Additional Background - The FCC Auction

There is a great deal of material describing the FCC auction design and the debate over it. An interested reader should consult McMillan [14] and the extensive material in the FCC's Docket No. 93.-253. The FCC adopted a series of general policies with respect to the kinds of auctions it would hold. These policies dictate different sales procedures for different situations. What follows is a general description of the process the FCC has adopted for the sale of the most valuable licences ("Broadband MTA").

These licences are being sold in a simultaneous progressive auction with discrete rounds. In these auctions, no sale of any licence takes place until the bidding is concluded on all licences. Bids are on individual licences are binding. (A bidder withdrawing a bid is still subject to covering the FCC's reduced revenue from the withdrawal, if any.) When there is a round of bidding with no bids that increase the price on any licence, each of the licences is sold to the bidder making the high bid on it. In order to keep bidders from holding back, there are "activity rules." Early in the auction, these are relatively lenient, but by the third and final phase of the auction, they are stringent. A bidder is considered active on a license in a round if he or she made the high bid for the license in the previous round or raises the bid on the license by at least the minimum amount required for the round. Bidders are allotted five "waivers" which can be used in an individual round to preserve their immediately prior activity status until the next round. Otherwise, in the first phase of the auction, bidders must remain active on licenses covering a population base which is at least one–third of the total population base for which they wish to remain eligible to bid. In the second phase, they must remain active on licenses covering at least two–thirds of their population base. In the final phase, they must remain active on all of it. That is, in the final phase, if a bidder ceases competing for a license for Philadelphia, for example, he can immediately switch to bidding on another license, but only if that other license covers a less populous region. The Broadband MTA auctions began with daily rounds but primarily proceeded at a two–rounds–per–day pace.

3 Additional Background - The Auction Design Problem

Only a few scholarly papers consider the auction design problems introduced by value interdependencies among different items for sale. Rothkopf [18] considered the bidder's problem in a simultaneous auction without combinatorial bids when there is an interdependency introduced by a constraint on the total of his or her bids. Smith and Rothkopf [21] consider the bidder's problem in a simultaneous auction without combinatorial bids when there is a fixed charge incurred if any bid succeeds. Rassenti, Smith and Bulfin [17], Banks, Ledyard and Porter [1] and McCabe, Rassenti and Smith [13] describe laboratory experiments with continuous–time auction mechanisms for simultaneous auctions with interdependent values. The first of these papers also states the integer programming problem that must be solved to select the set of nonconflicting combinatorial bids that maximizes revenue from the auction. Bykowsky, Cull and Ledyard [3] discuss "local equilibria" of simple and complex auction mechanisms. Krisha and Rosenthal [11] consider a simple model where "local" bidders compete for only one of two assets, while "global" bidders compete for both (without permitting combined bids), with a synergy should they win both. Aside from these papers and the material already cited about the design of the FCC's spectrum auction, we are not aware of any scholarly literature on this topic.

If frictionless aftermarkets existed for the assets being sold, the seller concerned only about efficient allocation would not have to worry about auction form. Almost any form would do since, whatever the initial allocation, the aftermarket would costlessly reallocate the assets efficiently. However, in the common situation in which there are significant transaction costs (and perhaps costly delays) associated with aftermarket transactions, the choice of auction form can affect economic efficiency. In addition, one of the important roles of auctions is to provide legitimacy by being demonstrably fair. The choice of mechanism may matter from that perspective as well.

The auction designer must choose between sequential sale and simultaneous sale, between a progressive process and one-time sealed bids, and between independent sales and allowing combinatorial bids. If the assets are sold independently, then bid evaluation presents no problem. If the assets are sold using one-time sealed bids, then there is a good deal of time available for determining the winning set of bids. Furthermore, the bid taker can resort to what we call a "political" solution of the bid selection problem–one that guarantees fairness and is likely to be rather effective as well. This "political" solution involves the bid taker finding the best feasible solution it can in a reasonable amount of time and then announcing it and giving all interested parties an opportunity to report a feasible solution with higher revenue. Apparent losing bidders will have incentive to explore bid combinations including their bids. If each of them is unable to find such a combination in a reasonable amount of time, no one is in a position to challenge the fairness of the bid acceptance process.

If the bidding is progressive, it can be either continuous or involve discrete rounds. If the bidding is continuous, then each new bid can be compared with high bids it must displace. If it exceeds all of them, then it becomes the current leading bid, and the other bids are displaced. This is a simple calculation. However, with combinatorial bids there is auction design choice. If unsuccessful bids are not kept available for use in the evaluation, the threshold problem can become quite serious. If they are, then there is a computational problem. In "AUSM", a computerized combinatorial auction procedure developed at the Cal Tech, this computational problem falls upon the bidders. They are given a list of unwithdrawn, currently unsuccessful bids. If they wish, they may incorporate one or more of the bids in this list into their bid.

An example might make this clearer. Suppose that there are four assets, A, B, C, and D for sale; that the current leading bids are \$100 for AB and \$200 for CD; and that there are unwithdrawn losing bids of \$50 for A and \$75 for D. A bidder for BC could combine these two losing bids with his bid of \$180 for BC to make a winning combination. The auctioneer would accept this bid because it increases the total revenue from \$300 to \$305. These bids would become the leading bids, and the formerly leading bids for AB and CD would, until they are withdrawn, be available to other bidders for making combinations of their own. In some variants of the Cal Tech procedure, bidders may submit unsuccessful bids in hopes that others may choose to use them in combination with their own bids.

If there are discrete rounds in the auction, as in the current radio spectrum auctions, then the computational problem of finding the winning bids must be solved by the bid taker at each stage. It must be solve much more rapidly, and the "political" solution described above is not available. In this context, the concern of the FCC about allowing combinatorial bids is understandable. However, the potential computational problems with allowing arbitrary combinatorial bids do not exist with certain structures of permitted combinatorial bids. The work presented in this paper is devoted to defining such structures and exploring their limits.

4 Formulation of the Problem

We begin by developing a mathematical formulation of the problem of selecting the revenue-maximizing set of bids.

Let X denote the set of all individual assets being auctioned. We assume that there are n assets, i.e. $|X| = n$. Any $B \subset X$ represents a combination of assets. Auction rules can specify which $B \subset X$ are **allowed combinations**, that is combinations of assets for which bidders may submit a bid. The family of all allowed combinations is denoted by $\mathcal{B}$, that is $\mathcal{B} = \{B \subset X : B \text{ is an allowed combination}\}$. Obviously, $|\mathcal{B}| \leq 2^{|X|}$.

In any outcome of a simultaneous auction, the winning combinations must be disjoint since no single asset can be sold to more than one bidder. More formally, an **outcome** is any $\mathcal{C} \subset \mathcal{B}$ such that $B \cap B' = \emptyset$ for every $B, B' \in \mathcal{C}$.

We concentrate on analyzing a single round of a simultaneous auction. Computational ease in determining winning bids in this case readily leads to a managable progressive auction by repeating the designed single–round model.

Throughout this paper we assume that the bidtaker's goal is to maximize revenue. (This might not always be the case. Sometimes bidtakers might want to establish the best political solution which is not necessarily the one that maximizes the revenue.)

Let $v_i(B)$ be the bid submitted by bidder i for the combination B. Let $w(B) = \max\{v_i(B) : i \text{ submitted a bid for } B\}$. If there is no bidder submitting a bid for B, we set $w(B) = 0$. In this notation, the bidtaker's goal is to find an **optimal outcome** $\mathcal{C}_{opt}$, i.e. to find the maximal amount he can get for the assets being auctioned:

$$\max_{\mathcal{C}} \sum_{B \in \mathcal{C}} w(B).$$

We also use notation $w(\mathcal{C}) := \sum_{B \in \mathcal{C}} w(B)$. In this notation, the problem is $\max_{\mathcal{C}} w(\mathcal{C})$.

Obviously, if B is in some outcome C, the bidtaker will sell all the assets from B to the bidder who submitted the highest bid for B. Therefore, the bidtaker will consider only $w(B)$ among all the bids submited for the allowed combination B.

We call any combination $B \in \mathcal{C}_{opt}$ a **winning combination**.

It is not hard to formulate the problem of finding an optimal outcome as 0-1 programming problem:

$$\max \sum_{B \in \mathcal{B}} w(B) x_B$$

with the constraints:

$$x_B \in \{0, 1\} \;\; \forall B \in \mathcal{B} \text{ and } x_B + x_{B'} \leq 1 \;\; \forall B \cap B' \neq \emptyset.$$

The problem of finding an optimal outcome is equivalent to a set–packing problem on a hypergraph $\mathcal{B}$ with weights $w(B)$ for every $B \in \mathcal{B}$. This problem is known to be NP-complete (Karp [10]). As we show later, the problem is easy for some classes of hypergraphs. Any such hypergraph leads to design of a managable simultaneous auction. We concentrate on hypergraph structures that (might) have applications in practice.

In our analysis of the complexity of finding an optimal outcome, we will use the standard asymptotic notation for growth of functions ($o(f(n))$, $O(f(n))$ and $\Omega(f(n))$). Precise definitions can be found, for example, in [4] or [7].

A related important issue in designing a simultaneous auction is that for every losing bid, i.e., combination B such that $B \in \mathcal{B} - \mathcal{C}_{OPT}$, there is a need for easy way to calculate the bid on B that would be necessary to make it a winning combination. In other words, keeping $w(B')$ unchanged for every $B' \in \mathcal{B}$, what is the minimum for $w_{\text{new}}(B)$ such that there exist some $\mathcal{C}_{OPT} \ni B$? In progressive auctions such information is crucial to every bidder. In order to find $w_{\text{new}}(B)$, we need to determine an optimal outcome, $\mathcal{C}^*_{OPT}$, for the set of assets $X^* := X - B$ and the set of allowable bids $\mathcal{B}^* := \{B' \in \mathcal{B} : B' \subset X^*\}$ with $w(B')$ as in the original auction.

Theorem 1 *Let $B \in \mathcal{B}$. B is an element of an optimal outcome for auction of X if and only if $w(B) = w(\mathcal{C}_{OPT}) - w(\mathcal{C}^*_{OPT})$ where $\mathcal{C}^*_{OPT}$ is as defined above.*

Proof: First, note that $w(B) \leq w(\mathcal{C}_{OPT}) - w(\mathcal{C}^*_{OPT})$. This is true because $\mathcal{C}^*_{OPT} \cup \{B\}$ is an outcome for an auction of X and $w(\mathcal{C}^*_{OPT} \cup \{B\}) = w(\mathcal{C}^*_{OPT}) + w(B) \leq w(\mathcal{C}_{OPT})$ since $\mathcal{C}_{OPT}$ is an optimal outcome.

If $w(B) = w(\mathcal{C}_{OPT}) - w(\mathcal{C}^*_{OPT})$, then $w(\mathcal{C}_{OPT}) = w(B) + w(\mathcal{C}^*_{OPT}) = w(\mathcal{C}_B)$ where $\mathcal{C}_B := \mathcal{C}^*_{OPT} \cup \{B\}$ is an outcome for auction of X. Since $w(\mathcal{C}_B) = w(\mathcal{C}_{OPT})$, $\mathcal{C}_B$ is an optimal outcome, and B is an element of an optimal outcome ($B \in \mathcal{C}_B$).

If $w(B) < w(\mathcal{C}_{OPT}) - w(\mathcal{C}^*_{OPT})$, then for any outcome $\mathcal{C} \ni B$ we can define an outcome $\mathcal{C}^* = \mathcal{C} - \{B\}$ for auction of X^*. Now, we have $w(\mathcal{C}) = w(\mathcal{C}^*) + w(B) \leq w(\mathcal{C}^*_{OPT}) + w(B) < w(\mathcal{C}_{OPT})$, which shows that $\mathcal{C}$ is never an optimal outcome. ▮

This theorem shows that determining the increment just sufficient to turn a losing combination into a winning one is as hard as determining an optimal outcome for the auction of X^*.

If the number of assets being auctioned is small, determining winning combinations is managable, but if the number of assets being auctioned is large, this problem for arbitrary $\mathcal{B}$ (e.g. $\mathcal{B} = 2^X$) will be unsolvable for practical purposes. However, for some special classes of $\mathcal{B}$ there exist fast and easy algorithms for finding an optimal outcome of the auction. Even in the most general cases, some combinations need not to be considered as candidates for winning combinations:

Observation 2 *Let $\mathcal{C}_{OPT}$ be an optimal outcome. Let $B_0 \supseteq B_1 \cup B_2 \cup \ldots \cup B_k$ where $B_0 \in \mathcal{C}_{OPT}$, $B_1, B_2, \ldots, B_k \in \mathcal{B}$ and $B_i \cap B_j = \emptyset$ for all $1 \leq i < j \leq k$ (B_i's are pairwise disjoint). Then $w(B_0) \geq \sum_{i=1}^{k} w(B_i)$.*

Proof: Define $\mathcal{C}' := (\mathcal{C}_{OPT} - \{B_0\}) \bigcup \{B_1, B_2, \ldots B_k\}$.

$$\sum_{B \in \mathcal{C}'} w(B) = \sum_{B \in \mathcal{C}_{OPT}} w(B) - w(B_0) + w(B_1) + w(B_2) + \ldots w(B_k).$$

Since $\sum_{B \in \mathcal{C}_{OPT}} w(B) = w(\mathcal{C}_{OPT}) \geq w(\mathcal{C}') = \sum_{B \in \mathcal{C}'} w(B)$ (because $\mathcal{C}_{OPT}$ is an optimal outcome), we get $w(B_1) + w(B_2) + \ldots w(B_k) \leq w(B_0)$ ▮

This observation shows that superadditivity (of the function w) plays an important role in determining winning bids. Combination B_0 will never be a winning combination if all the assets in B can be sold to other bidders for a larger total amount.

The following dynamic algorithm uses this idea to determine $\mathcal{C}_{OPT}$:

INPUT: $w(B)$ for all $B \subset X$ (if no bids are submitted or $B \notin \mathcal{B}$ then $w(B) = 0$).

1. For all $x \in X$, set $f(\{x\}) := w(\{x\})$, $C(\{x\}) := \{x\}$.
2. For $i = 2$ to n, do:
   For all $B \subset X$ such that $|B| = i$, do:
   (a) $f(B) := max\{f(B - B') + f(B') : 1 \leq |B'| \leq |B|/2\}$.
   (b) If $f(B) \geq w(B)$, then set $C(B) := B^*$ where B^* maximizes right hand side of (a).
   (c) If $f(B) < w(B)$, then set $f(B) := w(B)$ and $C(B) := B$.

3. Set $\mathcal{C}_{OPT} := \{X\}$.

4. For every $B \in \mathcal{C}_{OPT}$, do:
   If $C(B) \neq B$, then

   (a) Set $\mathcal{C}_{OPT} := (\mathcal{C}_{OPT} - \{B\}) \cup \{C(B), B - C(B)\}$.

   (b) Go to 4 and start it with the new $\mathcal{C}_{OPT}$.

After the first two steps $f(X)$ will be exactly $w(\mathcal{C}_{OPT})$ (i.e. $\sum_{B \in \mathcal{C}_{OPT}} w(B)$). $f(B)$ is calculated using the fact that all the assets from B will be sold to a single bidder (who submitted the highest bid for B) if and only if any dividing of B into smaller combinations will not lead to the higher value collected by the bidtaker. The variable $C(B)$ keeps track of the structure of $\mathcal{C}_{OPT}$ which is being determined recursively in the Step 4.

This algorithm requires fewer than $n^2 + \sum_{k=1}^{n} C 2^{k-1} \binom{n}{k}$ operations which is $\Omega(2^n)$ and $o(2^{2n})$. Obviously, the size of the input is the number of different combinations for which bids are submitted (we can set $w(S) := 0$ for any other $S \subset X$). It is possible that there will be a submitted bid for any $B \in \mathcal{B}$. In this limited sense, the algorithm is efficient if the set of allowed combinations $\mathcal{B}$ is large. For example, if bids for all $B \in \mathcal{B} = 2^X$ are submitted, the algorithm should be considered efficient since it is polynomial in the size of the input. However, if $n = 100$, the size of the input might be 2^{100} and if one billion submitted bids can be registered in a second, we would need more than $4 \cdot 10^{13}$ years to input all the data. The problem with this algorithm is that it requires a fixed number of steps regardless of the size of the input. Still, the algorithm might be useful for very small X.

As we just mentioned, there is no way of being sure that the simultaneous auction procedure is managable for large n if there are no restrictions on the set $\mathcal{B}$, not only because the problem is NP-complete, but also because the size of the input itself might become unmanagable.

In the presented algorithm, $f(X - B)$ is exactly $w(\mathcal{C}^*_{OPT})$ for every $B \subset X$ (a $\mathcal{C}^*_{OPT}$ can be determined by seting $\mathcal{C}_{OPT} := X - B$ and then executing Step 4.). In many cases, this information won't be a byproduct of an algorithm to find an optimal solution. However, the existence of an easy way to determine an optimal outcome automatically gives an easy way to determine an optimal outcome for an auction of the smaller set of assets $(X - B)$ provided allowed combinations on the smaller set of assets are of the same type as allowed combinations on the original set of assets.

Let us turn one last time to the presented algorithm. The algorithm doesn't need to evaluate the function f for every subset of X. It suffices to check only the sets $S \subset X$ in the smallest algebra $\mathcal{A} \subset 2^X$ of sets containing $\mathcal{B}$ (or only those $B \in \mathcal{B}$ for which a bid is submitted). An algebra $\mathcal{A}$ of sets is a family of sets $X \in \mathcal{A} \subset 2^X$ closed under union and taking complements, i.e. $A \in \mathcal{A} \Rightarrow A^c \in \mathcal{A}$ and $A, A' \in \mathcal{A} \Rightarrow A \cup A' \in \mathcal{A}$.

As will be demonstrated in the later sections, the managability of a simultaneous auction depends upon the structure of $\mathcal{B}$ rather than the size of $\mathcal{B}$ (the number of allowed combinations). Therefore, limiting the number of bids each bidder is allowed to submit, will not, in general, make the bid evaluation problem easier.

5 Nested Structures

The idea of evaluating bids for smaller combinations first and gradually evaluating bids for larger combinations gives an efficient algorithm for finding an optimal outcome if $\mathcal{B}$ forms a tree structure (a nested family of sets). A family of sets $\mathcal{B}$ forms a **tree structure** if for all $B, B' \in \mathcal{B}$, either $B \cap B' = \emptyset$ or B or B'. In other words, every two sets in the family are disjoint or one is a subset of the other. The sets in $\mathcal{B}$ define a directed tree $T(\mathcal{B})$ in a natural way: $\mathcal{B}$ is the set of vertices of $T(\mathcal{B})$, and (B, B') is an arc in $T(\mathcal{B})$ if and only if $B \supset B'$ and there is no $B'' \in \mathcal{B}$ such that $B \supset B'' \supset B'$. In other words, (B, B') is an arc in the tree $T(\mathcal{B})$ if and only if B covers B' in $\mathcal{B}$. Note that $B_1 \supset B_2$ if and only if there exist a directed path $P \subset T(\mathcal{B})$ from B_1 to B_2. If there is no directed path from B_1 to B_2 nor from B_2 to B_1 then $B_1 \cap B_2 = \emptyset$ since $\mathcal{B}$ forms a tree structure.

Note that we can always add X and all singletons to $\mathcal{B}$ without violating the tree structure property. In other words, if $\mathcal{B}$ forms a tree structure $\mathcal{B} \cup \{X\} \cup \{\{x\} : x \in X\}$ also forms a tree structure. Also note that every B has at most one ingoing arc (otherwise there would be two sets covering B and then their intersection would be at least $B \neq \emptyset$). If $X \in \mathcal{B}$, then every $B \neq X$ will have exactly one ingoing arc.

We call $B \in \mathcal{B}$ a **leaf**, if B has no outgoing arcs.

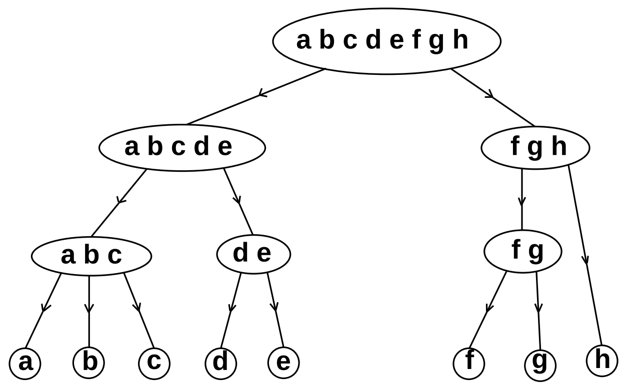

Figure 1: Example of a tree structure

For example, suppose that $X = \{a, b, c, d, e, f, g, h\}$ and allowable bids are all singletons and $\{a, b, c\}$, $\{d, e\}$, $\{a, b, c, d, e\}$, $\{f, g\}$, $\{f, g, h\}$, and X. Figure 1 shows $\mathcal{B}$ represented as a tree.

The following algorithm will produce $\mathcal{C}_{OPT}$:

INPUT: $T(\mathcal{B} \cup X)$ and $w(B)$ for all $B \in \mathcal{B}$.

1. Set $\mathcal{C}_{OPT}(B) = \{B\}$ for every leaf B.
2. For every $B \in \mathcal{B}$, calculate $d(B) :=$ the distance in $T(\mathcal{B})$ from X to B.
3. Find $B_{\max}$ such that $d(B_{\max}) \geq d(B')$ for all $B' \in \mathcal{B}$.
4. (a) Let B_0 be the tail of the unique ingoing arc of $B_{\max}$ (i.e. B_0 covers $B_{\max}$).
   (b) Let $\mathcal{S} := \{B \in \mathcal{B} : (B_0, B) \in A(T(\mathcal{B}))\}$. (Note that by the choice of $B_{\max}$ every $B \in \mathcal{S}$ is a leaf)
   (c) Calculate $w(\mathcal{S}) := \sum_{B \in \mathcal{S}} w(B)$.
   (d) If $w(B_0) > w(\mathcal{S})$, then set $\mathcal{C}_{OPT}(B_0) := \{B_0\}$.
   (e) If $w(B_0) \leq w(\mathcal{S})$, then set $w(B_0) := w(\mathcal{S})$ and $\mathcal{C}_{OPT}(B_0) := \bigcup_{B \in \mathcal{S}} \mathcal{C}_{OPT}(B)$.
5. If $B_0 = X$, then STOP ($\mathcal{C}_{OPT}(X)$ is an optimal outcome).
6. Set $\mathcal{B} := \mathcal{B} - \mathcal{S}$ and go to Step 3. (B_0 becomes a leaf in $T(\mathcal{B})$)

Step 1 requires $O(|\mathcal{B}|)$ time. Step 2 requires $O(|\mathcal{B}| \log(|\mathcal{B}|))$ time. Steps 3 and 4 require at most $O(n)$ time since any B can have at most $|B| \leq n$ outgoing arcs. These steps can be repeated at most $|\mathcal{B}|$ times since every time $\mathcal{B}$ is updated in Step 6 at least $B_{\max}$ is deleted from $\mathcal{B}$.

Theorem 3 *Let $\mathcal{B}$ form a tree structure. Then an optimal outcome can be determined in $O(n|\mathcal{B}|) \leq o(n^3)$ time.*

Proof: Let us first show by induction on $|\mathcal{B} \cup \{X\}|$ that $\mathcal{C}_{OPT}(X)$ from the algorithm is an optimal outcome. Obviously, the algorithm produces an optimal outcome for $\mathcal{B} \cup \{X\} = \{X\}$.

Let $\mathcal{S} := \{B \in \mathcal{B} : (X, B) \in A(T(\mathcal{B}))\}$. Note that $\mathcal{B}_B := \{B' \in \mathcal{B} : B' \subset B\}$ forms a tree structure and the algorithm produces an optimal outcome for every $\mathcal{B}_B$ where $B \in \mathcal{S}$. Also note that $\mathcal{C}_{OPT}(B)$ is this optimal outcome and that $w(\mathcal{C}_{OPT}(B)) = w(B)$ in the original algorithm. Since every $B' \in \mathcal{B}$ $(B' \neq X)$ intersects exactly one $B \in \mathcal{S}$, any optimal outcome for $\mathcal{B}$ is either $\mathcal{C}_{OPT} = \{X\}$ or a disjoint union of optimal outcomes in $\mathcal{B}_B$. Therefore, $\mathcal{C}_{OPT}(X)$ is an optimal outcome.

It remains to show that the algorithm requires at most $O(n \ln(|\mathcal{B}|) \leq o(n^3)$ time.

There can be at most $\lfloor n/k \rfloor$ sets $B \in \mathcal{B}$ with cardinality k (because they all must be disjoint). Therefore, $|\mathcal{B}| \leq \sum_{k=1}^{n} \lfloor n/k \rfloor \leq n \sum_{k=1}^{n} 1/k \leq n \ln n$. Therefore, the algorithm needs at most $O(|\mathcal{B}| \log(|\mathcal{B}|) + n|\mathcal{B}|) \leq O(n^2 \ln n) \leq o(n^3)$ time. ∎

Not only is it easy to determine an optimal outcome for a tree structures, it is also easy to determine how much higher the losing bid B_L must be to become a winning bid since $\mathcal{B}^* := \{B \in \mathcal{B} : B \subset X - B_L\}$ also forms a tree structure. Therefore, by Theorems 1 and 3, this can be determined in $o((n - |B_L|)^3) = o(n^3)$ operations.

6 Cardinality–Based Structures

A trivial simultaneous auction (which in fact does not allow combined bidding) is an auction where $\mathcal{B} = \{\{x\} : x \in X\}$. In such an auction, bids can be submitted for single assets only, and an optimal outcome is $\mathcal{C}_{OPT} = \mathcal{B}$ where the winning bids are the highest submitted bids for particular assets.

Even if arbitrary combined bids of size at most two are allowed, it is not at all trivial to determine the winning bids.

Observation 4 *Let $\mathcal{B} \subset \{B \subset X : |B| \leq 2\}$. Then finding an optimal outcome is equivalent to finding maximal weighted matching in a graph on at most $2n$ vertices.*

Proof: Let G' be a graph whose set of vertices is X and set of edges is the set of all combinations of size two, i.e. $\{x, y\}$ is an edge in the graph G' if and only if $\{x, y\} \in \mathcal{B}$. Let G be a graph obtained from G' by adding a vertex $\bar{x}$ and an edge $e_x := (x, \bar{x})$ with $w(e_x) := w(\{x\})$ for every $\{x\} \in \mathcal{B}$. Obviously, G has at most $2n$ vertices. Note that the obvious one to one correspondence between the edges of G and the combinations $B \in \mathcal{B}$ preserves disjointness, and therefore, $\mathcal{C}$ is an outcome if and only if corresponding edges are matching in G. ∎

Theorem 5 *(The maximum–weight matching algorithm) If $\mathcal{B} \subset \{B \subset X : |B| \leq 2\}$ then an optimal outcome can be determined in $O(n^3)$ time.*

Proof: The problem of finding $\mathcal{C}_{OPT}$ is equivalent to finding a maximum–weight matching in G which can be found in $O(n^3)$ time [8]. ∎

The weighted matching problem was first solved by Edmonds [6]. Algorithms for finding maximum–weight matching are not particularly transparent and are considered to be among the most complicated polynomial algorithms in combinatorial optimization. Therefore, the following result should not be suprising:

Theorem 6 *If $\mathcal{B} = \{B \subset X : |B| \leq 3\}$, then finding $\mathcal{C}_{OPT}$ is NP-complete.*

Proof: The problem of finding $\mathcal{C}_{OPT}$ is NP-complete in a particular special case of such an auction. Suppose that $w(B) = 0$ for all B such that $|B| \leq 2$ and $w(B) = 0$ or 1 for all other $B \in \mathcal{B}$. Then $\mathcal{C}_{OPT}$ is an optimal outcome if and only if $\mathcal{P} := \{B \in \mathcal{C}_{OPT} : w(B) = 1\}$ is an maximal 3-set packing on a hypergraph $\mathcal{B}$. The 3-set packing problem is known to be NP-complete ([7]). ∎

Therefore, allowing bids on any all combination of size three will lead to serious computational problems.

Another approach in designing a simultaneous auction is to allow only large combinations and singletons, i.e. $\mathcal{B} = \{B \subset X : |B| = 1 \text{ or } |B| > M\}$.

This might be appropriate, for example, where majority combinations give effective voting control, or where it is known that an efficient outcome implies some large combinatorial bid wins.

Theorem 7 *Let $\mathcal{B} = \{B \subset X : |B| = 1$ or $|B| > n/k\}$.*
Let $S := \{B \in \mathcal{B} : |B| > n/k$ and $w(B) > 0\}$.
Then $\mathcal{C}_{OPT}$ can be determined in $O(n|S|^k)$ time.

Proof: By Theorem 1, we can eliminate any $B \in S$ such that $w(B) \leq \sum_{x \in B} w(\{x\})$. This can be done in $O(|S|n)$ time.

Note that there can be at most k sets from S in any outcome $\mathcal{C}$. This is because every two sets in $\mathcal{C}$ must be disjoint and there can be at most k disjoint sets of size larger than n/k. Therefore, there are all together $\sum_{i=0}^{k} \binom{|S|}{i} = O(|S|^k)$ candidates for $\mathcal{C}_{OPT}$. For each such $\mathcal{C}$, we need to calculate $w(\mathcal{C})$ (This can be done in at most n steps.) and then find $\mathcal{C}_{OPT}$ among them. All together, we need $O(n|S|^k)$ time. ∎

$\mathcal{C}_{OPT}$ can be determined easily as long as the number of possible outcomes is not large. We need only to calculate $w(\mathcal{C})$ for every possible outcome $\mathcal{C}$. For example, if $k = 2$ in the Theorem, then we have at most $|S| + 1$ candidates for $\mathcal{C}_{OPT}$: every $B \in S$ and $B = \emptyset$ define an outcome:

$$\mathcal{C}_B := \{B\} \cup \{\{x\} : x \notin B\}.$$

This idea can be generalized for any measure of sets in X, not just the number of elements of the set.

When the importance of majority control motivates this kind of structure, then the use of alternate measures would be called for when different assets for sale are entitled to different voting weights.

Corollary 8 *Let μ be a finite measure on 2^X.*
Let $\mathcal{B} = \{B \subset X : |B| = 1$ or $\mu(B) > \mu(X)/k\}$.
Let $S := \{B \in \mathcal{B} : \mu(B) > \mu(X)/k$ and $w(B) > 0\}$.
Then $\mathcal{C}_{OPT}$ can be determined in $O(n|S|^k)$ time.

Proof: Note that there can be at most k sets from S in any possible outcome since μ is additive for disjoint sets. The rest of the proof is the same as in the proof of the Theorem. ∎

Without significantly increasing computational difficulty, combinations of size two can also be allowed. The following theorem compiles results from this section:

Theorem 9 *Let μ be a finite measure on 2^X.*
Let $\mathcal{B} = \{B \subset X : |B| \leq 2$ or $\mu(B) > \mu(X)/k\}$.
Let $S := \{B \in \mathcal{B} : \mu(B) > \mu(X)/k$ and $w(B) > 0\}$.
Then $\mathcal{C}_{OPT}$ can be determined in $O(n^3|S|^k)$ time.

Proof: As it was shown in the Theorem 7 and Corollary 8 there are only $O(|S|^k)$ different outcomes $\mathcal{C}$ containing disjoint sets from S (Maybe some additional sets from S can be eliminated using Observation 2 and $O(|X-B|^3)$ time is needed to check if B should be eliminated. All together, $O(|S|n^3)$ will suffice).

Let $\mathcal{C}_{OPT}$ be an optimal outcome for auction of $X^* = X - (\bigcup_{B \in S \cap \mathcal{C}} B)$ where $\mathcal{B}^* = \{B \subset X^* : B \in \mathcal{B}$ and $|B| \leq 2\}$. ${\mathcal{C}_{OPT}}^*$ can be determined in $O(|X^*|^3)$ time by Theorem 5. Then $\mathcal{C} \cup {\mathcal{C}_{OPT}}^*$ is a candidate for $\mathcal{C}_{OPT}$ with $w(\mathcal{C} \cup {\mathcal{C}_{OPT}}^*) = w({\mathcal{C}_{OPT}}^*) + \sum_{B \in S \cap \mathcal{C}} w(B)$.

Therefore, we can find $\mathcal{C}_{OPT}$ in $O(n^3|S|^k)$ time. ∎

As in the case of tree structures, determining how much higher the bid for the combination B_L must be in order to make B_L a winning combination is as hard as finding an optimal outcome for $X^* := X - B_L$ with $\mathcal{B}^* := \{B \in \mathcal{B} : B \subset X - B_L\}$. Obviously, $\mathcal{B}^*$ defines a cardinality based structure on X^* of the same type as $\mathcal{B}$ is on X.

7 Geometry–Based Structures

Often there exist some physical connection between the assets. For example, a group of neighboring assets may be more valuable as a combination than some otherwise equivalent random group of assets. An example of this type might be the sale of mineral rights on n tracts forming a single swath of offshore state waters. In such a situation most of the interesting combinations will be groups of neighboring assets, and this might lead to a design of a computationally managable combinatorial auction.

Let us suppose that there exists a total ordering among the assets. In other words, for every two assets $i, j \in X$: $i < j$ or $j < i$. For example, licences for radio frequencies in the cities on the East Coast can be ordered acording to North-South geographical ordering of the cities. Suppose that we allow combined bids for any set of consecutive assets $B_{i,j} := \{x \in X : i \leq x \leq j\}$ ($B_{i,j}$ is an **interval**). Then, as we show presently, $\mathcal{C}_{OPT}$ can be determined in $O(n^2)$ time. Note that if the assets can be put in the total order, then we can label them in a way that $X = \{1, 2, \ldots n\}$ with the ordinary relation $<$. We will use usual notation: $[n] := \{1, 2, \ldots n\}$.

If $X = [n]$ and $\mathcal{B} := \{B_{i,j} : 1 \leq i \leq j \leq n\}$, then the following algorithm will produce a $\mathcal{C}_{OPT}$:

INPUT: $w(B_{i,j})$ for all i, j.

1. Set $\mathcal{C}(1) := \{B_{1,1}\}$ and $W(1) := w(B_{1,1})$. Set $r := 2$.
2. Set $W(r) := w(B_{1,r})$ and $\mathcal{C}(r) := \{B_{1,r}\}$.
3. For $l = 2$ to r, do:
   If $W(l-1) + w(B_{l,r}) > W(r)$, then
   (a) Set $W(r) := W(l-1) + w(B_{l,r})$
   (b) Set $\mathcal{C}(r) := \mathcal{C}(l-1) \cup \{B_{l,r}\}$
4. If $r < n$, then set $r := r + 1$ and go to Step 2.
5. STOP ($\mathcal{C}(n)$ is an optimal outcome and $w(\mathcal{C}_{OPT}) = W(n)$).

Theorem 10 *Let $X = [n]$. Let $\mathcal{B} = \{B_{i,j} : 1 \leq i \leq j \leq n\}$ where $B_{i,j} := \{i, i+1, \ldots, j\}$. Then $\mathcal{C}_{OPT}$ can be determined in $O(n^2)$ time.*

Proof: Obviously, the presented algorithm needs $O(n^2)$ time. We will show by induction on n that $\mathcal{C}(n)$ is an optimal outcome. Clearly, $\mathcal{C}(1) = \{B_{1,1}\}$ is an optimal outcome when $n = 1$. Suppose that $\mathcal{C}(m)$ is an optimal outcome for the auction of first m assets whenever $m < n$. Let $\mathcal{C}_{OPT}$ be an optimal outcome for the auction of X. Let $\mathcal{C}^* := \mathcal{C}_{OPT} - \{B_{m+1,n}\}$ where $B_{m+1,n}$ is the unique set from $\mathcal{C}_{OPT}$ containing asset n. Note that $\mathcal{C}^*$ is an outcome for the auction of $[m]$. From Theorem 1 and the induction hypothesis, $w(\mathcal{C}^*) = W(m)$ and $w(\mathcal{C}_{OPT}) = W(m) + w(B_{m+1,n}) \leq W(n)$. The last inequality follows from the Step 3 of the algorithm when $l = m + 1$ and $r = n$. By optimality of $\mathcal{C}_{OPT}$, we conclude that the last inequality can be replaced by equality and $\mathcal{C}(n)$ is an optimal schedule. ∎

Note that for any $B_{i,j}$:

$$\mathcal{B}^* := \{B \in \mathcal{B} : B \subset X - B_{i,j}\} = \{B_{l,r} : 1 \leq l \leq r < i\} \cup \{B_{l,r} : j < l \leq r \leq n\}.$$

Therefore, determining how much higher the bid for $B_{i,j}$ must be to make $B_{i,j} \in \mathcal{C}_{OPT}$ amounts to determining $\mathcal{C}_{OPT}$ for two auctions of the same type as the original auction: auction of $[i-1]$ with $\mathcal{B}_{i-1} := \{B_{l,r} : 1 \leq l \leq r < i\}$ and auction of $\{j+1, \ldots n\}$ with $\mathcal{B}_{j+1} := \{B_{l,r} : j < l \leq r \leq n\}$.

Instead of considering intervals (sets of consecutive assets) on the line (totally ordered set) we can consider intervals on the circle (for example, offshore tracts around surrounding an island might define such a structure) and determine $\mathcal{C}_{OPT}$ by repeated use of the same algorithm. We apply the algorithm n times taking each of n elements to be the first one in the linear order defined by the cyclic order. For example, in k-th run of the algorithm we will consider the order: $k, k+1, \ldots, n, 1, 2, \ldots, k-1$. At the end, we compare n outcomes and choose $\mathcal{C}_{OPT}$ among them. More formally:

Corollary 11 *Let* $X = [n]$. *Let* $\mathcal{B} = \{B_{i,j} : 1 \leq i, j \leq n\}$ *where*
$B_{i,j} := \{i, i+1, \ldots, j\}$ *if* $i \leq j$ *and* $B_{i,j} := \{i, i+1, \ldots, n, 1, \ldots j\}$ *if* $i > j$.
Then $\mathcal{C}_{OPT}$ *can be determined in* $O(n^3)$ *time.*

Proof: Any outcome $\mathcal{C}$ contains at most one set $B_{i,j}$ where $i > j$ because 1 and n are in every such set and sets in $\mathcal{C}$ are disjoint. If an outcome $\mathcal{C}$ contains such $B_{i,j}$ then any other $B_{l,r} \in \mathcal{C}$ $j < l \leq r < i$ (because $\mathcal{C}$ is a collection of disjoint sets). If we rename all the assets $k < i$ into $k' := k + n$, then X becomes $\{i, i+1, \ldots, n, (n+1), \ldots, n+i-1\}$, $B_{i,j}$ becomes $B_{i,n+j}$ and any other $B_{l,r} \in \mathcal{C}$ becomes $B_{n+l,n+r}$. Therefore, $\mathcal{C}$ contains only sets satisfying conditions of the Theorem.

For every $1 \leq i \leq n$, the algorithm will determine an optimal outcome for the auction of $X_i := \{i, i+1, \ldots, n, (n+1), \ldots, n+i-1\}$ where $\mathcal{B} := \{B_{l,r} : i \leq l \leq r \leq n+i-1\}$. Since every outcome $\mathcal{C}$ is an outcome for the auction of X_i (Find $B_{i,j} \in \mathcal{C}$ that contains 1), it suffices to compare optimal outcomes for auctions of X_i. All this can be done in $O(n^3)$ time. ∎

In many cases, assets can be indentified as elements of a direct product of two linear orders. Geographic location of an asset (position on a map) is an example. Clearly, X can be represented as $X = [m] \times [n]$. The two dimensional analogues of the intervals are rectangles $B_{a,b;c,d} := \{(x,y) \in X : a \leq x \leq b\}$ and $\{c \leq y \leq d\}$. Unfortunately, there is no hope for finding a managable algorithm for an auction of X even if we allow only 2×2 rectangles to be allowable combined bids.

Theorem 12 *Let* $X = [n] \times [n]$. *Let* $\mathcal{B} = \{B_{a,a+1;b,b+1} : 1 \leq a < n$ *and* $1 \leq b < n\}$ *where* $B_{a,b;c,d} := \{(x,y) \in X : a \leq x \leq b$ *and* $c \leq y \leq d\}$. *Then finding a* $\mathcal{C}_{OPT}$ *is a NP-complete problem.*

Proof: As was shown in [2], the *optimal* 2×2 *salvage problem* is NP-complete. The input for this problem is an $n \times n$ grid and some set S of unit squares with integer coordinates. A 2×2 rectangle with integer vertices is called **functional** if all four of its unit squares belong to S. The problem is to determine a maximal number of functional non-overlapping (disjoint) 2×2 rectangles.

Any set S of unit squares with integer coordinates defines $w(B)$ for every $B \in \mathcal{B}$ in a natural way: set $w(B) := 1$ if B is functional and set $w(B) := 0$ if B is not functional. Then finding $\mathcal{C}_{OPT}$ is equivalent to solving the optimal 2×2 salvage problem. ∎

If we allow only rectangles of a specific type there may be computationally easy algorithm for finding $\mathcal{C}_{OPT}$. For example if we allow all the singletons and only $B := R_a := \{(a,y) \in X : y \in [n]\}$ for every $a \in [m]$ and $B := C_b := \{(x,b) \in X : x \in [m]\}$ for every $b \in [n]$ then $\mathcal{C}_{OPT}$ can be determined in a very transparent way. R_a and C_b can be viewed as rows and columns of a $m \times n$ rectangular grid.

This might be an appropriate representation of a situation in which a set of collectible assets is to be sold where assets have two different properties of interest to different collectors (for example, the year and the denomination of a coin).

$\mathcal{C}_{OPT}$ can be determined as follows:

INPUT: $w(B)$ for all $B \in \mathcal{B}$ (B's are singletons, rows or columns).

1. For every combined bid B (row or column), determine $w'(B) := \sum_{(x,y)\in B} w(\{(x,y)\})$.
   If $w'(B) \geq w(B)$, then set $w(B) := w'(B)$ (Note that in this case B can't be in any $\mathcal{C}_{OPT}$ by Observation 2.)
2. Calculate $w_R := \sum_{a\in[m]} w(R_a)$ (row sum)
   and $w_C := \sum_{b\in[n]} w(C_b)$ (column sum).
3. If $w_C \geq w_R$, then
   (a) Set $\mathcal{C}_{OPT} := \{C_b : b \in [n]\}$
   (b) For every $b \in [n]$, do: if $w'(C_b) = w(C_b)$, then
   $\mathcal{C}_{OPT} := \mathcal{C}_{OPT} \cup \{\{(x,b)\} : x \in [m]\} - \{C_b\}$ (replace C_b with singletons)
4. If $w_R > w_C$, then
   (a) Set $\mathcal{C}_{OPT} := \{R_a : a \in [m]\}$
   (b) For every $a \in [m]$, do: if $w'(R_a) = w(R_a)$, then
   $\mathcal{C}_{OPT} := \mathcal{C}_{OPT} \cup \{\{(a,y)\} : y \in [n]\} - \{R_a\}$ (replace R_a with singletons)

Note that in any outcome C there can't be a row and a column at the same time because $R_a \cap C_b = \{(a,b)\} \neq \emptyset$. Therefore, every outcome contains rows and singletons only or columns and singletons only. Since rows (columns) are pairwise disjoint, an optimal outcome containing rows (columns) is easily determined by checking whether replacing a row (column) by singletons increases $w(\mathcal{C})$. At the end, only the optimal outcome containing rows and an outcome containing columns need to be compared.

Example: Let X be vertices of a cube and let $\mathcal{B}$ consist of vertices (singletons) and edges (some doubletons) of the cube. Let $w(B) = 1$ for every vertex, let $w(e_1) = w(e_2) = w(e_3) = 5$, and let $w(e) = 3$ for every other edge (edges e_1, e_2 and e_3 are denoted in the Figure 2).

Note that there are three types of edges in the drawing: horizontal, vertical and "diagonal". If we were to follow the approach for the 2-dimensional case, we would find and evaluate optimal outcomes containing just one type of edges. In all three cases, we get $3+3+3+5 = 14$, and no singletons are involved. However, the optimal outcome $\{e_1, e_2, e_3, v_1, v_2\}$ has value $5+5+5+1+1 = 17$ and includes edges of different types. The reason why this approach doesn't work in this case is that there are edges of different types that are disjoint.

The following theorem generalizes this idea:

Theorem 13 *Let $\mathcal{B} = \mathcal{S} \cup \mathcal{F}_1 \cup \mathcal{F}_2 \cup \ldots \cup \mathcal{F}_k$ where $\mathcal{S} = \{\{x\} : x \in X\}$ (singletons) and where families $\mathcal{F}_i$, $1 \leq i \leq k$, are pairwise disjoint. Let $\bar{\mathcal{F}}_i := \mathcal{F}_i \cup \mathcal{S}$ and let $\mathcal{C}_{OPT}(i)$ be an optimal outcome for auction of X with allowed combinations $\bar{\mathcal{F}}_i$.*
If for every $B_i \in \mathcal{F}_i$ and $B_j \in \mathcal{F}_j$: $B_i \cap B_j \neq \emptyset$, then $\mathcal{C}_{OPT}$ is one of $\mathcal{C}_{OPT}(i)$, $1 \leq i \leq k$.

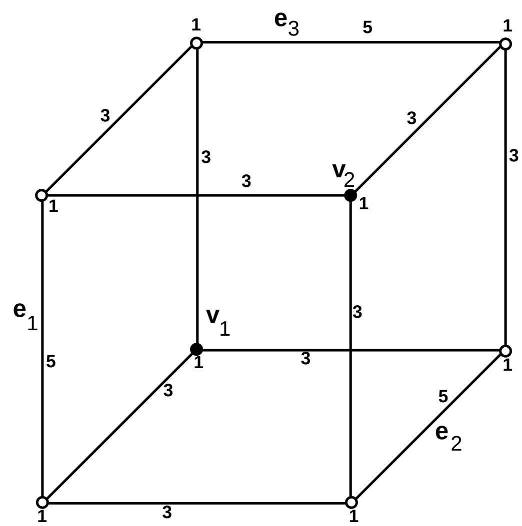

Figure 2: An example of auctioning the vertices and the edges of a cube showing that an optimal outcome can contain edges of all three orientations.

Proof: If $\mathcal{C}$ is an outcome, then it can't contain sets from more than one family $\mathcal{F}_i$ because every two sets from $\mathcal{C}$ must be disjoint. Therefore, every $\mathcal{C}$ is an outcome from $\bar{\mathcal{F}}_i$ for some i. Hence, it suffices to determine $\mathcal{C}_{OPT}(i)$ for every i and compare only those outcomes to determine an optimal one. ▮

Note that for any $B_L \in \mathcal{B}$, $\mathcal{B}^* = \{B \in \mathcal{B} : B \subset X - B_l\}$ still satisfies conditions of the theorem and determining how much a losing bid must be raised is as hard as finding an optimal outcome for $X - B_l$ with allowed combinations from $\mathcal{B}^*$. However, it is easier to find $w(\mathcal{C}_{OPT}(i)^*)$ where $\mathcal{C}_{OPT}(i)^*$ is an optimal outcome for $\bar{\mathcal{F}_i}^* := \{B \in \bar{\mathcal{F}}_i : B \subset X - B_L\}$ where $B_L \in \bar{\mathcal{F}}_i$. $w_{new}(B_L)$ should be larger than $w(\mathcal{C}_{OPT}) - w(\mathcal{C}_{OPT}(i)^*)$.

The 2-dimensional example satisfies the conditions of the theorem with $\mathcal{F}_1$ being the set of all rows and $\mathcal{F}_2$ being the set of all columns. Both $\bar{\mathcal{F}}_1$ and $\bar{\mathcal{F}}_2$ form a tree structure where $d(X, B_{\max}) = 2$.

We can generalize this to the d-dimensional case. Let $X = [n_1] \times \ldots \times [n_d]$. Let $H_a^i := \{(x_1, \ldots, x_d) \in X : x_i = a\}$ (a hyperrow in the ith coordinate). Let $\mathcal{F}_i := \{H_a^i : 1 \leq a \leq n_i\}$. Then the conditions of the theorem are satisfied, and it is easy to determine $\mathcal{C}_{OPT}(i)$ since each $\bar{\mathcal{F}}_i$ forms a tree structure with $d(X, B_{\max}) = 2$. For example, if X is a cube then $\mathcal{B}$ consists of vertices (singletons) and faces (hyperrows) of the cube.

Obviously, the $\mathcal{F}_i$ don't need to have the same structure:

Example: Let $X = [n_1] \times [n_2] \times [n_3]$. Let $\mathcal{T} \subset 2^{[n_1]}$ be a tree structure on $[n_1]$, let $\mathcal{D} \subset 2^{[n_2]}$ be a collection of singletons and doubletons, and let $\mathcal{L} \subset 2^{[n_3]}$ be a collection of intervals on the totally ordered set $[n_3]$. Let

$$\mathcal{F}_1 := \{\cup_{a \in I} H_a^1 : I \subset \mathcal{T}\},$$

$$\mathcal{F}_2 := \{\cup_{a \in I} H_a^2 : I \subset \mathcal{D}\},$$

$$\mathcal{F}_3 := \{\cup_{a \in I} H_a^3 : I \subset \mathcal{L}\}.$$

Note that $\mathcal{T}$, $\mathcal{D}$, and $\mathcal{L}$ are actually families of index sets. Conditions of the theorem are satisfied, and we can find an optimal outcome $\mathcal{C}_{OPT}(i)^*$ for each of the $\mathcal{F}_i$ by Theorems 3, 5 and 10 if we consider hyperrows as singletons. In order to find $\mathcal{C}_{OPT}(i)$, it suffices to check if any of the hyperrows from $\mathcal{C}_{OPT}(i)^*$ can be replaced by singletons. After necessary replacements are done, $\mathcal{C}_{OPT}$ can be determined by comparing $\mathcal{C}_{OPT}(i)$, $i = 1, 2, 3$. Althought the structure of $\mathcal{B}$ is rather complicated, $\mathcal{C}_{OPT}$ can be determined in $O((n_1)^3 + (n_2)^3 + (n_3)^2)$ time.

8 Concluding Remarks

Some auction markets sell a considerable number of assets simultaneously. Often, the value of an asset to a bidder depends upon which other assets he also wins. In such situations, allowing bidding on combinations of assets may offer a way of increasing the efficiency of the allocation of assets, but it can raise computational problems. Those who determine the rules for simultaneous auctions must determine for which combinations to allow bids. In sales with many assets, "all combinations" may not be a workable answer, and "no combinations" may not be a desirable one. In this paper, we have considered restricted sets of combinations for which combinatorial bidding presents a provably manageable computational burden. We hope that this will allow auction designers to design workable auctions that are more efficient.

Deciding for which combinations to allow bids puts a responsibility on auction designers—a responsibility that some of them may find politically risky to fulfill. We note that computational impossibility does not provide a protection against such responsibility, and that "no combinations" is only one choice among what we have now established are many feasible possibilities. Responsible auction designers will try to determine the kinds of combinations of greatest economic significance and attempt to allow bids upon them in the auction—at least when there is reason to believe that economies of scale exist. The politically astute among them may well try to involve the potential bidders in that determination process. Perhaps, the Department of the Interior's use of a nomination process in deciding which offshore tracts to offer for sale at a given time could serve as a model.

Acknowledgement

The authors thank Richard Engelbrecht-Wiggans for helpful comments.

References

[1] Banks, J., J. O. Ledyard and D. Porter, "Allocating Uncertain and Unresponsive Resources: An Experimental Approach", *RAND Journal of Economics* **20** (1989), pp. 1-23.

[2] Berman, F., D. Johnson, T. Leighton, P. W. Shor, and L. Snyder, "Generalized Planar Matching", *Journal of Algorithms* **11** (1990), pp. 153-184.

[3] Bykowsky, M. M., R. J. Cull and J. O. Ledyard, "Mutually Destructive bidding: The FCC Auction Design Problem", *Social Science Working Paper, California Institute of Technology* **916** (1995).

[4] Cormen, T. C., C. E. Leiserson and R. L. Rivest, Introduction to algorithms, *MIT Press, Cambridge, MA*, 1990.

[5] Federal Communications Commission, "In the Matter of Implementation of Section 309(j) of Communication Act, Competitive Bidding: Notice of Proposed Rule Making," PP Docket No. **93-253** (1993), Released October 12.

[6] Edmonds, J., "Maximum Matching and a Polyhedron with 0,1 Vertices," *J. Res. NBS* **69B** (1965), pp 125-130.

[7] Garey, M.R. and D.S. Johnson, Computers and intractability : a guide to the theory of NP-completeness, *Freeman, S.F.*, 1979.

[8] Lawler, E.L., Combinatorial Optimization: Networks and Matroids *Holt, Rinehart and Winston, N.Y.*, 1976.

[9] Lovasz, L. and M.D. Plummer, Matching Theory *North-Holand, N.Y.*, 1986.

[10] Karp, R.M., "Reducibility among Combinatorial Problems," Complexity of Computer Computations, R.E. Miller and J.W. Thatcher (eds.), *Plenum Press, NY* (1972), pp. 85-103.

[11] Krishna, V. and R. Rosenthal, "Auctions with Synergistic Values," Economics Dept. working paper, Boston University, (1995).

[12] McAfee, R. Preston, 1993, "Auction Design for Personal Communictions Services," PacTel Exhibit in PP Docket No. **93-253**.

[13] McCabe, K.S., S. Rassenti and V.L. Smith, "Smart Computer Assisted Markets", *Science* **254**, pp. 534-538.

[14] McMillan, John, "Selling Spectrum Rights," *Journal of Economic Perspectives*, **Vol 8** (1994), pp. 145-162.

[15] McCurdy R. C., Private Communication (1981), November 3.

[16] Milgrom, Paul R. and Robert B. Wilson, "Affidavit of Paul R. Milgrom and Robert B. Wilson," Before the Federal Communications Commission, PP. Docket No. **93-253** (1993), November 8.

[17] Rassenti, S. J. , V. L. Smith, and R. L. Bulfin, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation", *Bell Journal of Economics* **13** (1982), pp. 402- 417.

[18] Rothkopf, Michael H., "Bidding in Simultaneous Auctions with a Constraint on Exposure", *Operations Research* **25**(1977), pp. 620-629.

[19] Rothkopf, Michael H., "Bidding Theory: The Phenomena to be Modeled," Auctions, Bidding and Contracting: Uses and Theory, Richard Engelbrecht-Wiggans, Martin Shubik and Robert Stark, Eds., *New York University Press, New York* (1983), pp. 105- 120.

[20] Schackmann, A., Private Communication (1989).

[21] Smith, Stephen A., and Michael H. Rothkopf, "Simultaneous Bidding with a Fixed Charge if Any Bid Succeeds", *Operations Research* **33** (1985), pp. 28-37.