Chapter 16.

Non-computational Approaches to Mitigating Computational Problems in Combinatorial Auctions

by

Aleksandar Pekeč, The Fuqua School of Business, Duke University

and

Michael H. Rothkopf, School of Business and RUTCOR, Rutgers University

16.1 Introduction

Hikers encountering a fallen tree blocking a trail can climb over it, cut a path through it, or walk around it. In general, obstacles can be overcome, reduced, or avoided. Often, reducing the obstacle or avoiding it is a preferable choice. The computational problems in combinatorial auctions are no different. The previous four chapters have been largely devoted to describing ways of overcoming them. While it is important to have the ability to overcome unavoidable computational difficulties, good combinatorial auction design will certainly want to take advantage of appropriate ways to reduce or avoid them. That is the topic of this chapter.

This section of this chapter sets the context by reviewing briefly the computational issues in combinatorial auction design, the context of auction design including the information available to the designer, and properties that the auction designer must trade off in selecting the auction format and procedures. The following four sections then look at mitigation opportunities prior to bid submission, at the time of bid submission, after bid submission but prior to announcing a tentative set of winning bids, and after the announcement of a tentative set of winning bids.[1] A final section concludes with a discussion of the implications of these opportunities for auction design.

16.1.1. Computational issues

A distinguishing feature of combinatorial auctions is the computational complexity of tasks considered trivial and mundane in non-combinatorial auctions. The central problem that has gathered considerable attention in recent years is that of winner determination to which first three chapters of part III of this book are devoted. The winner determination problem is discussed in detail in Lehmann *et al.* (Chapter 12), as well as in Müller (Chapter 13) and in Sandholm (Chapter 14). Here, we briefly note that the winner determination problem in its simplest form is equivalent to the set packing problem (as pointed out by Rothkopf *et al.* 1998).[2] Thus, the winner determination problem is one of the basic NP-complete problems and the fundamental source of potential computational difficulties in combinatorial auctions.

In this chapter, we focus on combinatorial bids that are "package bids" only, that is we will assume that all bids are simple "all or nothing" bids. Our observations readily

---

[1] There is no section on using aftermarkets to avoid computational problems. If perfect aftermarkets were to exist, there would be no need for combinatorial auctions. With imperfect aftermarkets, bidders' values in the combinatorial auction should reflect the opportunities that do exist in these aftermarkets. Given the results of a combinatorial auction, even imperfect aftermarkets may present opportunities to mitigate the effects of misallocations in the auction. To the extent that this can happen, it can tip the balance in the combinatorial auction design tradeoff between efficient allocation and lower transaction costs towards lower transaction costs. Note also that aftermarkets can not only deal with misallocations caused by the auction design, but also by misallocations caused by changed circumstances and by bidder errors.

[2] We refer the reader to Chapters 12-14 to learn more about variations of the winner determination problem.

generalize to the more general case of Boolean combinations of "package bids" (defined as "combinatorial bids" in this book).[3]

While it is fundamental, the winner determination problem is not the only issue that could bring computational difficulties. Among other problems that could be of computational concern is calculating how much more a losing bidder must bid in order for his bid to become a winning one, the problem of resolving ties, and determining minimum bid increments in an iterative auction. However, most of these problems, regardless of the approach, are inherently related to the winner determination problem and its complexity. For example, methods for determining bid increments have to revolve around calculations of the gap (see glossary) for each losing bid. As shown in Rothkopf et al. (1998), the complexity of calculating gap for any losing bid is equivalent to the complexity of the winner determination problem.[4] Similarly, resolving ties by random selection (e.g., FCC, 2002) is essentially equivalent to solving a winner determination problem.[5] Thus, most of the potentially difficult computational problems in combinatorial auctions are equivalent or closely related to the winner determination problem. Hence, without serious loss of generality, one can focus on addressing complexities of the winner determination problem.

---

[3] The "all-or-nothing" feature that can be represented by the AND Boolean operation, and not other Boolean operators such as OR and XOR, is a fundamental generator of complexities. Suppose Boolean combinations using OR clauses of XOR clauses of single items are allowed (but AND clauses that could create "packages" are not allowed). Then "winner determination" becomes almost trivial computationally as the problem reduces to the max weight system of distinct representatives, i.e., max weight bipartite matching, which is an optimization problem solvable in polynomial time.

[4] Given the computational demands of calculating gap for all losing bids in a combinatorial auction in which solving the winner determination problem is not easy, several approximate approaches have been developed (e.g., see Hoffman *et al.*, Chapter 17 and Kwasnica *et al.*, 2004). For example, Hoffman (2001), reported that in the simulations of one of the earlier versions of the FCC's combinatorial auction, more than 99% of the computational time was spent on calculating the gap for all losing bids; this problem eventually resulted in the change of the auction rules (FCC, 2002).

[5] There are other approaches to resolving ties. Several such approaches are discussed in Pekeč and Rothkopf (2003).

16.1.2 Context

When designing an auction and dealing with computational complexity, an auction designer has to take into account the importance of various potentially desirable properties of an auction and, if necessary, make appropriate tradeoffs. Some of desirable properties are:[6]

- allocative efficiency, i.e., maximizing the total value to the winners of the items being auctioned;
- revenue maximization (or payment minimization);
- low transaction costs and auction speed as both the bidders and the bid taker care about their costs of participating in the auction;
- fairness, that is, concern about equal treatment of competitors (and the appearance of it);
- failure freeness, as failures should be minimized and their impact mitigated;
- transparency;
- scalability is important in design of auctions that will be used repeatedly.

An important concern in government auctions doesn't have to be one in commercial procurement auctions, and vice versa. For example, a government auction might have to pay special attention to allocative efficiency, while cost minimization could be a primary goal in corporate procurement; a government might not be in position to appear unfair or afford settling for a suboptimal allocation (as it could face lengthy lawsuits), while corporate auctions could aim at speeding up the auction procedure at the price of possibly

[6] Following Pekeč and Rothkopf (2003)

failing to find the optimal allocation. Regardless of the goals of the particular situations, some non-computational auction design approaches discussed in this chapter could reduce the complexity burden while preserving (most of) the other desirable properties.

Another concern to the auction designer is the potential informational burden. One aspect of such a burden is mere handling information that could be massive given that there are up to 2^m-1 potential packages of m items that could be bid on. Clearly, eliciting bidder valuations for all possible packages (or all possible allocations) could be even more informationally demanding.[7] A different aspect of informational burden arises when the auction designer wants to or has to release comprehensive but aggregated information about the auction, such as the gap for all bids or minimum bid increments for all biddable combinations in every round of an iterative auction.[8] Therefore, auction design choices on the auction format, including the information flow from bidders to the bid-taker and vice versa could affect the implementability of the auction from the information management point of view. Some possibilities at the disposal of an auction designer are discussed in this chapter.

16.2. Mitigating Complexity Prior to Bid Submission

There are many tasks that need to be done before an auction begins. This section discusses how these tasks can mitigate computational complexity. The possibilities discussed below include the choice of definition of the items to be sold, the definition of a metric to make items comparable, and the definition of which combinations are to be biddable.

---

[7] See Segal (Chapter 11).

[8] In fact, as reported by Hoffman (2001), the burden of calculating precisely the gap for all bids can be overwhelming in the context of the FCC combinatorial auction.

Sometimes the definition of an "item" is obvious. Often, however, there is considerable discretion involved in defining the items to be auctioned. For example, when the U.S. Federal Communications Commission decides to sell the rights to use 30 MHz of radio spectrum, it could offer one 30 MHz license, two 15 MHz licenses, a 20 MHz license and a 10 MHz license, three 10 MHz licenses, or 30 1 MHz licenses. It can also divide the licenses geographically, and sell five regional licenses, 50 state licenses, or 500 local licenses. Clearly, the way the assets to be sold are divided up into items can have a profound effect on the computational problems of winner determination. An artless choice, for example selling 15,000 local 1 MHz licenses when economic use of the spectrum requires regional aggregations and at least 10 MHz, could lead unnecessarily to horrendous computational difficulties. Sellers need to use their knowledge of the desires and the economic situation of their potential buyers to make sensible definitions of items. Sellers normally have much such knowledge and potential buyers often have incentive to provide or improve it. Indeed, a seller who does not have and cannot get any knowledge of what his potential buyers want should probably sell his assets competitively in one lot to the potential resellers who best know the market.

There is little that can be said in general about the decision about how to divide assets for sale into lots. However, even though it depends upon the particulars of the situation, it is a critical auction design issue. There are some possibilities that are of potential value that are worth mentioning.

One of these is defining some sort of measure that can be used, perhaps with adjustments, to make items fungible. When this can be done, it greatly simplifies the computational problem of selecting winners. An example of this is daily electricity supply auctions. In the best of these, a megawatt of power injected into the electricity grid over a given hour at point A is equivalent to a megawatt injected into the grid that hour at point B except for an adjustment designed to take account appropriately of the costs of transmission congestion. This is called "Locational Marginal Pricing." Electricity auctions that ignored this congestion effect and assumed that congestion-free zones could be predefined have run into difficulties. Recently, there has been a proposal to auction both energy and transmission rights simultaneously.[9]

Another one is to predefine biddable combinations in a way that reflects the underlying economics and will simultaneously mitigate potential computational difficulties during the course of an auction. For example, bids for take off and landing slots at an airport could be limited to pairs involving one landing slot and one take off slot. This would meet the economic need of airlines to take off any plane that they land at an airport while leading to a computationally tractable auction.[10] The next three subsections discuss these possibilities in more detail.

16.2.1. Defining items for sale

The issue of defining items for sale is not specific to combinatorial auctions. For example, what has to be sold and how should it be sold depends on physical nature of the items for sale, (dis)synergetic valuations that the bidders have over potential items that could result in either bundling or unbundling of such items, as well as on the practical considerations of conducting an auction in a timely and efficient manner. In many

---

[9] For that proposal, see O'Neill *et al.* (2002). For earlier discussions of alternative definitions of transmission rights see Hogan (1992), Chao and Peck (1996), Baldick and Kahn (1997) and Chao *et al.* (2000).

[10] See Rothkopf *et al.* (1998). Also note that one of the first papers on combinatorial auctions, Rassenti *et al*. (1982) focuses on this potential application. See also Ball *et al.* (Chapter 20).

situations, the basic physical units that cannot reasonably be further decomposed are clear. For example, well-running used cars should not normally be sold as used car parts. However, there are many other important situations where such atoms do not exist such as radio frequencies, electricity contracts, and undeveloped land division. As for possible synergies and dissynergies, if these occur naturally or are common to all bidders, a good auction design will define auction items by appropriately bundling or unbundling in order to accommodate such situations. Thus, trying to bundle objects that are synergetic to every bidder into an auction item, and unbundling them into separate auction items when dissynergies are common, could be a general rule of thumb. However, one has to consider possible budget constraints and differences in individual bidder valuations since in such situations forced bundling might be inefficient and revenue deficient.[11] Similarly, even if bundling some objects is not synergetic, such "inefficient" bundling could make the total number of items to be auctioned small and the auction process fast and manageable. For example, an auctioneer selling off the property of a bankrupt restaurant might well chose to lump together into a single lot all of the pots and pans rather than sell them separately reasoning that any possible lost revenue will be more than offset by savings in transaction costs.

In general, bidders might not agree on synergies and dissynergies. Thus, perfect, noncontroversial bundling or unbundling in defining auction items might be difficult. One approach to such situations would be for all bidders to agree on the items that should

---

[11] For example consider two objects *a* and *b*. where bidder 1 has values $v_1(a)=4, v_1(b)=4, v_1(ab)=9$; bidder 2 has $v_2(a)=1, v_2(b)=5, v_2(ab)=7$, bidder *3* has $v_3(a)=5, v_3(b)=1, v_3(ab)=7$. Suppose also that each bidder has a budget constraint of *8*. If bidding on *ab* only is allowed, bidder 1 could win with bid $7+\varepsilon$, while if bids on unbundled objects a and b are allowed, bidders 2 and 3 will win with bids of $4+\varepsilon$ on each of *a* and *b*.

not be further decomposed, and then to allow bids on any combination of such items.[12] This approach could be useful for defining items when there is no obvious physical description or limitation. The drawback is that bidders might have infinite number of objects they would consider buying[13] and, even in the finite case, the resulting number of items could be unmanageably huge. Thus, perfect, noncontroversial item definitions may not be possible or practical.

However items to be auctioned end up being defined, an auction designer has to decide whether to allow bids on combinations of items (provided that there is more than one item). As many chapters of this book discuss (e.g., Nisan Chapter 9, Segal Chapter 11, Lehmann *et al.* Chapter 12, and Leyton-Brown *et al.* Chapter 19) and as briefly discussed in the introduction, combinatorial bidding could introduce complexities and potentially insurmountable obstacles in conducting an auction.

If many combinatorial bids are placed in an auction, this could mean (among other things) that bidders have serious conflicting synergetic valuations and/or that some items could have been bundled before being auctioned off. Taking this observation to the extreme, the final pricing and allocation, if done prior to the auction, could eliminate the need for the auction itself and lead to optimal posted prices since it would eliminate any interest for bidding except exactly for the winning bidders submitting the winning bids. So, in addition to allocation and price discovery, combinatorial auctions are mechanisms for optimal bundling discovery. The process of discovering the optimal bundling of items is the one that differentiates combinatorial from non-combinatorial auctions, and

[12] These items would be atoms of the algebra generated by the union of all objects (each object is a set) that any of the bidders might be interested in bidding on. For example, if each bidder provides a finite list of frequency ranges they would consider buying, one can define auction items to be all non-empty intersections of any collection of such frequency ranges (possibly from different bidders).

[13] For example, a bidder could be interested in any 20MHz range in 600-700 MHz band.

the one that is responsible for inherent complexities of combinatorial auctions. Thus, to the extent possible, auction designers should aim toward understanding likely optimal bundles. This understanding could be more critical to the auction success than understanding likely auction prices and likely auction winners. In turn, this suggests that an auction designer should put an effort into properly defining auction items in order to manage implementation complexities of combinatorial bidding.[14]

For example, consider an auction of three items *a, b, c,* with highest valuations (all from different bidders) as follows: *v(a)=3, v(b)=4, v(c)=5, v(ab)=10, v(bc)=10, v(ac)=10, v(abc)=13*, and with second highest valuations being exactly one less (and placed by a completely different set of bidders). Then, by posting the following prices, *p(a)=p(b)=p(c)=5, p(ab)=10, p(bc)=p(ac)=11, p(abc)=14*, the seller can bundle and allocate the items optimally by simple posted price mechanism, instead of running the combinatorial auction. Of course, this assumes that seller has guessed these prices correctly or that he has some information on optimal bundles and prices, i.e., the type of information that is usually discovered by the auction mechanism. Note that knowing high valuations without the knowledge of specific structure is not sufficient to find an optimal allocation; all two-item bundles are valued at *10* but the optimal pricing has to price-discriminate in order to clear the market in an optimal allocation that requires bundling *a* and *b*. Also note that, with the assumed high valuations, the seller could simply decide to accept only bundled bids on *ab* and bids on *c*, and not lose anything by limiting bidding in this way. Thus, discovering optimal bundling is valuable for market-mechanisms that allow combinatorial bids.

[14] As will be discussed in subsection 16.2.3, not all combinatorial bids are equally cumbersome during the course of an auction. The informational burden on the auctioneer is also a factor (as discussed in 16.1.2).

It is worth mentioning that even when packages can be divided among bidders, how this is done can have a major impact on the effectiveness of the auction. The original California day-ahead electricity auctions purchased 24 commodities, electricity in each of the hours of the next day. No bids on combinations of hours were allowed, and fractions of bids could be accepted. This is an awkward definition for a potential bidder with a generating plant that has start up costs, minimum run levels, and requires four hours to start or stop. A better design proposed by Elmaghraby and Oren (1999) would have bidders offering electricity from hour A to hour B. Note that the acceptance of a fraction of a bid in this auction would affect the allocation of start up costs but not the feasibility of starting and stopping the generating plant. An even better approach, actually used by the New York system operator and the operator of the Pennsylvania-New-Jersey-Maryland system considers bids involving start up costs and minimum run levels, solves a mixed integer program to find the optimum feasible dispatch for the day, uses hourly market clearing prices based on variable costs and then further compensates bidders whose fixed costs would not be covered so that they will not lose money at the optimum dispatch.[15]

In summary, the way auction items are defined has direct impact on the level of complexities of a combinatorial auction. Even if items seem to be naturally defined, choosing to bundle or unbundle some of such items could significantly aid the process of optimal bundle discovery. Thus, if combinatorial bidding is to be allowed and if the computational complexity of running an auction is an issue, defining auction items

[15] That such equilibrium market-clearing prices always exist to support the MIP solution is shown in O'Neill *et al*. (2004). The rules for the day-ahead and real-time energy auctions (as well as those for other products, such as operating reserves) are in an almost constant state of refinement. For the latest rules in these electricity markets, see www.pjm.com and www.nyiso.com.

should be looked at as an opportunity to mitigate effectively the potential for encountering computational nightmares while running the auction.

16.2.2. Defining units

Bids on different combinations are usually incomparable,[16] and this incomparability is one of the core issues in combinatorial auctions. One potential way to simplify this problem is to define an underlying measure on all possible biddable combinations and use this measure when solving the winner determination problem and other computational tasks such as determining minimum bid increments in multi-round action formats.

For example, in the latest FCC combinatorial auction, the MHzPop measure, i.e., the bandwidth of the spectrum multiplied by the population covered by the geographical scope of the license, is used in calculating minimum bid increments. The bid value is divided by the MhzPop value of the underlying combination and in this way all bids can be compared on the one-dimensional scale of $/MhzPop. Electricity auctions use $/MWh (dollars per megawatt hour)—adjusted by location for transmission congestion in the better auction designs and within predefined zones in others—as a common measure.[17]

In fact, many iterative combinatorial auction proposals suggest use of some underlying measure to define minimum bid increments.[18] The use of an underlying measure can be exact or approximate. If it is exact, then the items are truly fungible, as

[16] Suppose $b(ab)=5$ and $b(bc)=7$. While, in isolation, the bid on *ab* is smaller than the bid on *bc*, in the presence of other bids, say $b(a)=2$ and $b(c)=5$, one could argue that the bid on *ab* is "better" than bid on *bc*.
[17] There is a substantial literature on this. See, for example Chao *et al.* (2000), Baldick and Kahn (1997), and especially O'Neill *et al.* (2002).
[18] For example, the RAD mechanism described in Kwasnica *et al.* (2004) and the FCC combinatorial auction mechanism (FCC, 2002).

with treasury bonds, and individual items need not be differentiated at all. Use of an approximate simplification in determining minimum bid increments might not be viewed as too problematic, especially in iterative mechanisms where there are multiple chances to remedy the effects of any approximations. However, even in iterative mechanisms, there may be exceptions such as when auction activity is coming to a close and the simplification proposes a minimum bid increment that is too high and that eliminates potential bidders who could drive up the price.

The one-dimensionality that results from introducing an underlying measure to compare bids, can simplify the winner determination problem, too.[19] For example, there is a proposal (Rothkopf and Bazelon 2003) for auctioning residual spectrum subject to the limited rights of existing licensees. Since the existing rights holder has a strangle hold on the use of the residual rights, there will tend to be only one bidder for a given license's residual rights. The proposal is to take bids on many such licenses but sell only those getting the highest bid per MhzPop. Since spectrum of high frequencies is considerably less valuable per MhzPop than spectrum of lower frequencies that can penetrate buildings, the proposal is to auction low and high frequencies separately.

The key to introducing bidding units through an underlying measure on biddable combinations is that such measure is widely accepted and appropriate for the situation in hand as with KWhs adjusted for transmission congestion and MhzPops for spectrum with sufficiently similar frequencies and that everyone involved in the process, especially bidders and the auctioneer, is aware of the potential to reach a suboptimal outcome. In that sense, choice of the underlying measure that is aligned with bidders' values is critical

[19] Some relevant theoretical analysis is Müller (Chapter 13).

for success of such approach. Thus, situations in which a measure that is acceptable to all bidders exists, are prime candidates for this approach.

16.2.3. Defining biddable combinations.

Allowing bids on prespecified combinations of items can mitigate the difficulty of the winner determination problem, as well as that of other potentially computationally intractable issues that have to be resolved during the course of a combinatorial auction. Müller (Chapter 13) is devoted to structures of biddable combinations that ensure computational tractability for the winner determination problem.[20] Two things are worth noting. First, in most situations, it is not the size of the biddable combinations nor their number, but rather their structural properties (how they intersect, complement and supplement each other) that is the main determinant of the complexity of the winner determination problem. Second, most other computational problems in combinatorial auctions involve solving some sort of a winner determination problem, so focusing on complexity of the winner determination is of primary importance.

An important concern in limiting biddable combinations is that such limits could give an unfair advantage to bidders who are able to express their synergies using biddable combinations over bidders whose synergies lie across combinations that are not biddable. This again points out to the importance of properly defining auction items and of understanding bidders' synergies, as careful choices there could allow for restricting combinatorial bids in a way that won't be perceived as *unduly* limiting or unfair. Keep in mind that if computational complexity is an issue, there may be no set of usable rules that

[20] This is a well-defined and general combinatorial optimization question. However, there is little practical reason to dwell on the analysis of structures that have little chance of being relevant to any auction. Richer payoffs will be found in the exploration of structures of potential practical use.

is completely neutral. In particular, allowing no combinations at all may greatly favor some bidders over others.

It is also worth noting that decisions on limiting biddable combinations interact strongly with decisions on defining items. It may be fairer and more efficient to have more items with bidding allowed on a limited but well chosen set of combinations than to lump the items into a few biddable "super-items" and allow bidding on any possible combination of these super-items.

16.3 Mitigating complexity during bid submission

One way for the bid taker to deal with the computational complexities of a combinatorial auction is to request that bidders submit bids together with information that will help in the computational process. Two specific ideas of this type are discussed in this section. One approach is to completely shift the computational burden from the auctioneer to the bidders. In the other, the auctioneer is completely responsible for computation but allows bidders to guide the process; in this way, a heuristic can find a suboptimal solution that is aligned with bidders' preferences that are expressed at the time of bid submission.

16.3.1 Relegating complexity to bidders

The auctioneer could choose to take a passive role and let bidders not only submit their bids, but also prove to the auctioneer that their bids ought to be winning ones. Some of the very first and successful combinatorial auction designs, such as AUSM (Banks *et al.,* 1989) have this feature. Examples of such policies include standard auction designs that relegate computational burden to bidders such as AUSM and PAUSE (see Land *et al.*, Chapter 6). The general idea here is that the auctioneer expects bidders to present a collection of (non-intersecting) bids that improves on the current best collection. (Usually, the measure is the revenue for the auctioneer.) Thus, it is the bidders who have to solve the winner determination problem in such designs. There are variants of such procedures that might generate better auction results. Bidders could be allowed to submit bids without having to demonstrate that their bid could be combined with other bids into the best available collection of bids. In this way, even bidders without any computational abilities could participate in an auction. Then, the auctioneer (and perhaps other entities, e.g., those who have an interest in particular rivals not winning) could compute the best collection among available bids.

Relegating computational burden to the bidders is an option that does not really mitigate the computational complexities of combinatorial auctions, but it does relieve the auctioneer at the expense of participating bidders. The advantage of such a scheme is that the bidders know which combinations are of interest to them, while the auctioneer may be less well informed and have to be prepared to consider combinations that will not be bid. When designing such auction procedures, one has to be careful about the possible burden of managing in timely manner what may turn out to be massive amounts of information. Also, this approach assumes continuous bidding or multiple rounds of bidding as bidders have to be aware of all bids currently in the system when composing and submitting their proposals.

16.3.2 Bidder prioritization of combinations

Park and Rothkopf (2004) proposed a fundamentally different approach. They suggest that the bidders be allowed to bid on as many combinations of whatever kind they like but that they be required to prioritize their combinations. They propose that the bid-taker evaluate the bids, starting with no combinations, then including each bidder's top priority combination, then including each bidder's two top priority combinations, etc. until either all combination have been included or the time for computation has expired. This approach takes advantage of the bidders' knowledge of which combinations are important. It assures a measure of fairness when computational considerations do not allow all desired combinations to be considered, and it takes advantage of the fact that integer programming algorithms often perform much better in practice than worst-case based bounds. Note that the Park and Rothkopf approach need not be limited to giving equal number of allowable combinations to each bidder. It could be generalized to accommodate any prespecified construction of the lists of combinations to be considered based on the bidders' input preferences lists. For example, in an iterative auction, bidders who are more active could be favored by being allowed to have more combinations considered.

A potential concern with this method, as well as with any other limited search method (as discussed in section 16.4.1) is that bidders might have incentives to submit bids with a primary goal of complicating computational process in order to limit search and perhaps influence the final, possibly suboptimal, allocation that favors them.

16.4 Mitigating complexity prior to allocation decisions

Sandholm (Chapter 14) discussed methods that do not guarantee finding an optimal solution to the winner determination problem or do not guarantee finding it in a reasonable time-frame. The previous section discussed how the problem of solving the winner determination problem can be relegated to the bidders or how solving it could be guided based on bidder input. This section focuses on possibilities of mitigating complexities after bids are submitted.

16.4.1. Limited search

A simple general method to deal with the complexity of solving winner determination problem or any other complex problem during the course of an auction is for the bid-taker to announce (prior to the auction start) an upper bound in terms of computational time and/or other resources to be devoted to solving any particular problem. For example, in an iterative combinatorial auction, the time for solving the winner determination problem between two rounds could be limited, and the best available solution when the time expires could determine the provisional winners.

This approach is almost uniformly used in determining minimum bid increments in iterative combinatorial auctions. Instead of basing the increment value on the value of the gap (which would involve solving as many winner determination problems as there are biddable combinations), one could abandon any computational effort and prescribe a simple fixed increment amount (e.g., as in one of the first commercially implemented procurement combinatorial auctions; see Ledyard *et al.*, 2002.) or could use the linear programming relaxation of the winner determination problem in order to provide an approximation.[21]

Resorting to a limited search option opens up several issues:

- Should bidders know the details of the algorithm used to solve the winner determination problem (or another problem) under time/resource constraint and should bidders be able to replicate the procedure used by the bid-taker?

- The very fact that a suboptimal solution could be selected potentially allows for a new gaming dimension. At least in theory, bidders might submit bids aimed at slowing the algorithm down and potentially bringing it to settle for a suboptimal solution that favors them. However, in order to implement such strategy in any limited search approach to the winner determination problem, one would have to know intricate details of the winner determination problem computation, have a very good idea about high bids on all relevant biddable combinations, and have considerable computing resources and expertise, likely beyond those available to the bid-taker. Thus, except maybe in specific narrow situations,[22] it is likely that such malicious bidding with a primary goal of hindering computation would not be a serious problem. The following two observations explain why one might so conclude: If a bid is close to being includable in an optimal solution, then its maker risks it being accepted. Hence, it would be risky to make it if it were insincere. However, if it is not close to being acceptable, it may well be eliminated easily in any branch and bound calculation of the winning bids. Furthermore, outsourcing computation as suggested in section 16.4.2

[21] See Hoffman et al. (Chapter 17).

[22] For example, government auctions with a dominant corporate bidder, could exhibit all such properties: government might have to be very open about the computational process while at the same time bidders could have overwhelming resources and have good ideas of rivals' valuations.

could eliminate any computational advantage a bidder has over the bid taker making such an approach even more problematical.

- Given limited resources, the auction designer can choose a heuristic or an approximation algorithm that doesn't guarantee finding an optimal solution but does have some performance guarantees (e.g., that the solution found is not too far in some measure from the optimal one). If such an algorithm finds its solution before all resources are exhausted, e.g. before time expires, should the remaining resources be used to attempt to improve on the proposed solution (if it is not provably optimal)?

- What if several problems, e.g., the winner determination problem and the minimum bid increment problem, have to be solved within joint resource constraint? How should resources be allocated?[23]

- How should complaints of bidders who would be winners in an optimal solution but are not winners in the limited search solution, be handled?

16.4.2. Outsourcing search

A compromise approach between relegating complexity to bidders and the limited search option is to allow but not require bidders and, perhaps, other parties to participate in the computational effort. This approach aims at separating computation from the allocation decision and was first proposed by Pekeč (2001). An auction mechanism could treat computation as a transaction service that the auctioneer can outsource. For

[23] For example, suppose that the winner determination problem is solved by a suboptimal solution and, based on that solution, minimum bid increments get calculated optimally without exhausting all resources. Suppose further that not enough resources are left to recalculate the minimum bid increments in case the leftover resources are used for improving the solution to winner determination problem. Should this further improvement be attempted?)

example, the auctioneer could find an initial solution of the winner determination problem and invite everyone (bidders and perhaps others) to submit alternative solutions. Note that this approach requires the public release of bids (but not necessarily identification of the bidders). In some contexts, as when bidders' bids could reveal sensitive commercial information, this might be a disadvantage. (See Rothkopf *et al.* 1990)

There are several incentives issues that might complicate implementation of this approach. First, the auction-designer could try to set some incentives for computation contributors. Perhaps, the auctioneer could pay a fee to whoever finds the best solution to the winner determination problem by, e.g., awarding the party that first submits the best solution (hopefully, a provably optimal one) some combination of a fixed monetary award and a percentage of the improvement that the submitted solution made relative to the initial one. Second, as discussed above participants interested in winning items in the auction might have incentives to submit bids that are aimed at complicating computation, while at the same time having no chance to become winning bids. Similarly, some concern might arise with respect to those interested in collecting computation fees.

In summary, limiting resources for computation surely brushes away potential computational disasters, but it does raise incentive issues.

16.5. Mitigating complexity after preliminary allocation decisions

A key goal of auctions is often fairness. Perfect efficiency, while desirable, is unlikely to be achieved in large, complicated combinatorial auctions. Indeed, efficiency may well be traded off against the transaction costs associated with conducting the auction. While perfect efficiency may be unattainable, good efficiency and fairness can be obtained even if it proves impossible to get a provably optimal solution to the winner determination problem. One way to do this is the "political" solution suggested in Rothkopf *et al.* (1998). The essence of such a political solution to the winner determination problem is to give an opportunity to bidders (and, perhaps, other parties) to challenge and improve upon a proposed allocation before it is made final. Not only will providing such an opportunity for challenges provide an opportunity to improve the solution to the winner determination problem, it assures fairness. The essence of this is that it will be impossible for a bidder to challenge suboptimal auction results as unfair if the bidder himself has had a fair opportunity to suggest an improved solution. The reason simple auctions are deemed fair is that a bidder who had a fair opportunity to bid cannot credibly complain about the price received by a rival whom he failed to outbid. Similarly, a bidder who has a fair opportunity to improve upon a proposed, but possibly suboptimal solution to the winner determination problem cannot credibly complain if, later, a better solution is found. Since he had a fair chance to find it, it was clearly too difficult to find in the time and with the optimization technology available.

It is worth noting that this "political" approach to solving the winner determination problem is likely to be highly effective. It can be thought of as a decomposition approach in which each bidder is assigned the task of finding a way to have some of his tentatively excluded bids included in the final allocation. If the values at stake matter, each of the bidders with tentatively excluded bids will be highly motivated. Bidders, of course, are free to retain optimization experts as auction consultants just as they now hire economists.

While allowing bidders the chance to improve upon a potentially suboptimal solution to the winner determination problem will assure fairness, it is not necessary to limit the parties who may suggest improvements to bidders. In particular, it may at times make sense for the bid taker to allow any party to suggest improvements. As discussed above, it could motivate such participation by offering whoever supplies the best solution a portion of the improvement in the objective function achieved. This can be thought of as an "economic" solution to the winner determination problem.

16.6 Conclusions

Combinatorial auction design, like many other design problems, is an art. Its practitioners must make choices that affect conflicting design objectives, and its results must be evaluated in the context of the facts of the particular situation. In this chapter, we have attempted to describe a variety of ways that combinatorial auction designers can achieve good auction results when computational issues are potentially intractable. There are quite a few possibilities, and some of them have attractive features of potential use in important contexts.

The items that serve as the underlying atoms of the auction can (and need to) be defined artfully so as to make computation and other aspect of the auction workable. Where bid takers have knowledge of bidders' preferences, and they normally will, this needs to be taken into account in defining the items. These preferences also need to be taken into account in deciding which combinations will be biddable, and the item definition decision and the biddable combination decision need to be made together. When possible, one should take advantage of ways of making items fungible, either exactly or approximately.

If computation is a problem, its burden can be left with the auctioneer or shifted to the bidders during the course of the auction. If the auctioneer retains it, the bidders may be asked to prioritize bids on combinations so that if the computation cannot consider all of the combinations, it will have considered all of the most important ones. The auctioneer can outsource post-bidding computation to computational experts or to the bidders themselves. Further, the auctioneer can prevent potential problems that could arise from suboptimal allocations by allowing for challenges by third parties.

The combination of the possibilities discussed in this chapter with the computational capabilities discussed in many of the others will allow much better designs for combinatorial auctions in a wide variety of challenging contexts.

References

Baldick, R. and E. Kahn 1997. Contract Paths, Phase Shifters, and Efficient Electricity Trade. *IEEE Transactions on Power Systems* **12**(2) 749-755.

Banks J.S., J.O. Ledyard, D. Porter 1989. Allocating uncertain and unresponsive resources: An experimental approach. *RAND J. Econ.* **20** 1-25.

Chao, H.-P., S. Peck, 1996. A Market Mechanism for Electric Power Transmission. J. Regulatory Econ. **10**(1) 25-59.

Chao, H.-P., S. Peck, S. Oren, and R. Wilson 2000. Flow-based Transmission Rights and Congestion Management. *The Electricity Journal* **13**(8) 38-58.

Elmaghraby, W. and S. Oren 1999. The Efficiency of Multi-Unit Electricity Auctions. *The Energy Journal* **20**(4) 89-116.

FCC 2002. The Federal Communications Commission Public Notice DA02-260, available at http://wireless.fcc.gov/auctions/31/releases.html#da020260.

Hobbs, B.F., M.H. Rothkopf, L.C. Hyde, R.P. O'Neill 2000. Evaluation of a Truthful Revelation Auction for Energy Markets with Nonconcave Benefits. *J. of Reg. Econ.* **18**(1) 5-32.

Hoffman, K. 2001. Issues in Scaling Up the 700 MHz Auction Design. Presented at *2001 Combinatorial Bidding Conference*, Wye River Conf. Center, Queenstown, MD, October 27.

Hogan, W. W., 1992. Contract Networks for Electric Power Transmission. J. Regulatory Econ. **4**, 211-242.

Kwasnica, A. M., J. O. Ledyard, D. Porter, and C. DeMartini, 2004, "A New and Improved Design for Multi-Object Iterative Auctions," *Management Science*, to appear.

Ledyard, J.O., M. Olson, D. Porter, J.A. Swanson, D.P. Torma 2002. The First Use of a Combined Value Auction for Transportation Services. *Interfaces* **32**(5), 4-12.

O'Neill, R. P., U. Helman, B. F. Hobbs, W. R. Stewart Jr., and M. H. Rothkopf 2002. A Joint Energy and Transmission Rights Auction: Proposal and Properties. *IEEE Transactions on Power Systems* **17** 1058-1067.

O'Neill, R. P., P.l M. Sotkiewicz, B. F. Hobbs, M. H. Rothkopf, and W. R. Stewart, Jr., 2004, "Efficient Market-Clearing Prices in Markets with Nonconvexities," to appear in the *European Journal of Operational Research*.

Park, S., M.H. Rothkopf 2004. Auctions with Endogenously Determined Allowable Combinations. RUTCOR Research Report 3-2001, Rutgers University., to appear in the *European Journal of Operational Research*.

Pekeč, A. 2001. Tradeoffs in Combinatorial Auction Design. Presented at *2001 Combinatorial Bidding Conference*, Wye River Conf. Center, Queenstown, MD, October 27.

Pekeč, A., M.H. Rothkopf 2000. Making the FCC's First Combinatorial Auction Work Well. An official filing with the Federal Communications Commission, June 7, 2000, available at: http://wireless.fcc.gov/auctions/31/releases/workwell.pdf

Pekeč, A., M.H. Rothkopf 2003. Combinatorial Auction Design. *Management Science* **49** 1485-1503.

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

Rothkopf, M.H., C. Bazelon 2003. Spectrum Deregulation without Confiscation or Giveaways. *New America Foundation*, Washington, DC.

Rothkopf, M.H., A. Pekeč, R.M. Harstad 1998. Computationally Manageable Combinational Auctions. *Management Science* **44** 1131-1147.

Rothkopf, M.H, T.J. Teisberg, E. P. Kahn 1990.Why Are Vickrey Auctions Rare? *Journal of Political Economy* **98** 94-109.