MICHIGAN
ROSS SCHOOL OF BUSINESS

Working Paper

Simplified Bidding and Solution Methodology for Truckload Procurement and Other VCG Combinatorial Auctions

Damian R. Beil
Stephen M. Ross School of Business
at the University of Michigan

Amy Cohn
University of Michigan
College of Engineering

Amitabh Sinha
Stephen M. Ross School of Business
at the University of Michigan

Ross School of Business Working Paper Series
Working Paper No. 1034

Original Date: May 2006
This Revision: March 2007

UNIVERSITY OF MICHIGAN

Simplified Bidding and Solution Methodology for Truckload Procurement and Other VCG Combinatorial Auctions

Damian R. Beil, Amy Cohn, Amitabh Sinha
University of Michigan, Ann Arbor, MI 48109
dbeil@umich.edu, amycohn@umich.edu, amitabh@umich.edu

March 2007*

Abstract

In theory, combinatorial auctions can provide significant benefits in many real-world applications, such as truckload procurement. In practice, however, the use of such auctions has been greatly limited by the need for bidders to bid on an exponential number of bundles and for the auctioneer to solve an exponentially large winner-determination problem. We address these challenges for VCG combinatorial procurement auctions in which a bidder's cost for each bundle is determined by a cost function with an amenable structure. For example, the cost to a trucking company of servicing a bundle of loads is based on the least-cost set of tours covering all of these loads, which can be found by solving a simple minimum cost flow problem. Leveraging the fact that true-cost bidding is a dominant strategy in VCG auctions, we suggest that the bidders' challenges can be overcome by specifying this true-cost function explicitly as a bid, rather than computing and communicating each bid individually. Moreover, we propose to embed this true-cost function directly within the winner-determination problem, using the strength of mathematical programming to solve this problem without ever explicitly enumerating the bids. The research challenge is then to identify this cost function, and formulate and solve the corresponding winner-determination problem. We focus primarily on how this can be done for the truckload procurement problem, outline a more general framework for the approach, and identify a number of other promising applications.

Keywords: Combinatorial Auction; Procurement; Mathematical Programming; Truthful Bidding

*This is a revised version; original May 2006

1 Introduction

The average U.S. manufacturer spends 60% of its revenue to procure goods and services (Bureau of the Census 2005). Internet auctions are increasingly used as a procurement method, enabling faster negotiations with broader sets of potential suppliers. Auctions can also be used to execute complex negotiations which would be impossible with traditional methods. Often such auctions are combinatorial in nature, enabling suppliers to express economies of scale and scope via their bids and preferences.

Combinatorial auctions have been successfully used in practice for billions of dollars worth of transactions. Perhaps the most famous is the series of Federal Communications Commission auctions (Cramton 1997, 2002, Binmore and Klemperer 2002, Klemperer 2002). Other applications of combinatorial auctions include transportation services (at Sears Logistics, as documented by Ledyard et al. (2002)), airline landing slots at airports (Rassenti et al. 1982), operations procurement at GE (General Electric Corporation 2001), and strategic sourcing at Procter & Gamble (Sandholm et al. 2006). Truckload procurement has been among the most popular applications of such combinatorial auctions (Caplice and Sheffi 2006, Song and Regan 2003, Sandholm et al. 2006).

Nevertheless, two major hurdles prevent the full realization of the benefits of combinatorial auctions. The first is communication-based: To completely express economies of scale and scope among all items being auctioned, bidders must construct and submit bids on an exponential number of subsets of items (called bundles). For example, in a truckload procurement auction, hundreds or thousands of loads are auctioned simultaneously, which results in a huge number of possible combinations of loads to be considered. The second hurdle is computational: The auctioneer must solve a winner determination problem over the corresponding exponential number of bids, which in general is highly intractable. The implications of this are pointed out, for instance, by de Vries and Vohra (2003) and Pekec and Rothkopf (2003). Although there has been much recent research into the amelioration of these challenges, which we survey in §2, they continue to present a significant obstacle to the practical use of combinatorial auctions.

Both hurdles stem from an underlying assumption that all bids must be explicitly enumerated and communicated, which is true in the most general case Nisan and Segal (2006), for example, in which costs are exogenously endowed. However, in many situations the bids themselves can naturally be seen to arise from some clearly defined function. In this paper we suggest that rather than enumerating the specific bid for each of the exponentially many bundles, the bidder instead simply communicate the bid-generating function to the auctioneer. This addresses the bidder's hurdle. The auctioneer, on the other hand, is still potentially burdened with the task of computing all bids and solving the resulting exponentially large winner determination problem. Our second insight is that the auctioneer can embed the bid-generating function directly into the winner determination problem. The resulting auction is outcome-equivalent to the explicitly enumerative auction, and under certain conditions permits solution which under the explicitly enumerative auction would be intractable.

Surprisingly little work takes advantage of a bid-generating function explicitly in both communicating bids and solving the winner determination problem. One exception we are aware of is Hobbs et al. (2000), who explicitly use the cost function in bidding in an energy market auction, but do not directly address how the winner determination problem is solved. (We review related combinatorial auctions literature in the next section.) Most work to tackle the hurdles described above use iterative auction or bidding language approaches, which both generally take specific bids as primitives and seek to communicate them to the auctioneer. In contrast, we bypass communicating specific bids and take the bid-generating function as the primitive. Iterative auctions can be shown in some settings to converge to optimality without requiring bidders to bid exhaustively on all bundles (Sandholm and Boutilier 2006), and results such as polynomial communications for bidding languages can be shown for specific preference structures (Nisan 2006). Boutilier (2002) empirically demonstrated that the winner determination problem can be solved much faster using a logic-based bidding language which concisely expresses the bidders' underlying bid structure, compared to solving the auction when all combinatorial bids are explicitly specified. Similarly, our approach is most appropriate if the bid-generating functions can be succinctly communicated to the auctioneer and the winner determination problem which embeds these functions is tractable. In general, no known solution to the hurdles described above applies for all possible problem types and bidder preference structures (Nisan and Segal (2006) proved that with a fully general preference structure, it is impossible to find the optimal allocation without exponential communication in the worst case). However, in this paper we show that a rich and important group of real-world problems can be tackled using mathematical programming techniques to formulate and solve the bid-generating function and the corresponding winner determination problem.

For instance, consider the truckload procurement problem: An auctioneer wishes to procure truckload services to transport a set of loads, each load consisting of a full truckload and specified by an origin and a destination. In this case, the cost to a trucking company of serving a bundle of loads is based on the least-cost set of tours covering all of these loads, in recognition of the fact that drivers must ultimately return to their home bases. This cost can be computed using a simple minimum cost flow problem, which constitutes the cost-generating function. If the auctioneer conducts a VCG auction (the auction mechanism that is the focus of this paper, described in detail in §2.1), then the cost-generating function is in fact the bid-generating function, because truthful bidding is a dominant strategy in VCG auctions. Therefore, in our proposed methodology, the bidders would simply transmit the parameters of these bid-generating functions (i.e., the parameters of the minimum cost flow problems of each bidder), rather than explicitly listing their bids for each possible bundle of loads, which would be the case in a canonical VCG combinatorial auction. Resultantly, the auctioneer's problem is simply a multicommodity flow problem, rather than a set partitioning problem.

1.1 Our Contributions: Embedding Bid-generating Functions in Combinatorial Auctions

Our main contribution is to identify the benefits of using bid-generating functions directly within the bidding and winner determination of VCG combinatorial auctions. In particular, our methodology leverages application-specific structures within a mathematical programming framework to tractably achieve auction outcomes equivalent to a fully-enumerated combinatorial auction. In this paper, we

- introduce the notion of bidding via bid-generating functions (rather than computing and enumerating an exponential number of explicit bids) and, more importantly, embedding the bid-generating function itself (rather than the specific bids) directly into the mathematical programming formulation of the winner-determination problem; and
- explore in detail the application of this idea in truckload procurement auctions, providing a framework for bidding and winner determination in this context.

In addition, we prove that this approach is outcome-equivalent to a VCG auction in which all bids are explicitly enumerated; we identify a number of other real-world applications for which the underlying bid-generating function may enable a similar approach; and we outline several new areas of research that stem from the fundamental ideas presented here.

1.2 Paper Outline

We begin with a review of the related literature in the next section. In §2.1, we provide a brief primer into VCG auctions, which is the auction mechanism used in this paper. Our exploration of truckload procurement auctions appears in §3. We subsequently generalize our methodology in §4. In §5, we discuss a few other application areas and how our methodology may be applied in these areas. We conclude in §6 with a brief discussion of our work as well as the future research questions it generates.

2 Literature Review

Auctions have been used for millennia to leverage competition and find market clearing prices; the reader interested in general auction theory is referred to a recent text by Krishna

(2002). In this paper, our focus is on combinatorial procurement auctions. Combinatorial auctions were surveyed by de Vries and Vohra (2003) as well as Pekec and Rothkopf (2003), while a recent book edited by Cramton et al. (2006) provides a comprehensive examination of the theory and applications of combinatorial auctions in various domains.

Applications of combinatorial auctions in practice are diverse, as mentioned in the second paragraph of §1. These applications have spawned long-term contracts for truckload shipments (Caplice and Sheffi 2006, Sheffi 2004, Song and Regan 2003), airline landing slot allocation (Ball et al. 2006, Rassenti et al. 1982), wireless spectra (McMillan 1994, Cramton 1997, 2002, Binmore and Klemperer 2002), material sourcing at consumer-goods companies (Sandholm et al. 2006), and even school lunch programs (Epstein et al. 2002). These applications have spawned an entire industry specializing in combinatorial auctions, with some of the bigger names being CombineNet and Manhattan Associates. This area is also related to "smart markets", as studied by McCabe et al. (1991), and the more general field of electronic markets (Anandalingam et al. 2005, Wu and Kleindorfer 2005).

In this paper, we focus on procurement (or reverse) auctions. While in theory procurement auctions are identical to the forward auctions where items are sold, they nonetheless hold specific contextual challenges and have received a great deal of attention. In addition to the procurement papers mentioned earlier, for example, Chen et al. (2005) study multi-unit auctions for supply chains, Che (1993) studies multi-dimensional auctions, and Hohner et al. (2003) study the combinatorial auctions for strategic goods conducted by Mars, Inc.

Despite the widespread use of combinatorial auctions, the problems we alluded to earlier due to the combinatorial explosion in the number of bundles remain. A number of approaches have been suggested in the literature to tackle this issue. One stream of research restricts the bidding language, so that bidders may only bid on a subset of the potential bundles (Nisan 2000). An example of this is the XOR-of-OR bidding language, which was used by the FCC in their 2000 spectrum auction (Günlük et al. 2005). Another approach is preference elicitation, where the auctioneer proactively asks bidders to submit bids for specific bundles, with the aim of quickly uncovering a "good" set of bundles (Sandholm and Boutilier 2006). Many auction mechanisms (including the ones cited above) for combinatorial auctions are *iterative*, where bidders submit bids in multiple rounds. Since bidders submit only a small number of bids in each round, the combinatorial explosion can be overcome in those cases where the iterations converge. Parkes (2006) is a recent survey of iterative combinatorial auctions, with some other work including that of Kwon et al. (2005) and Ausubel et al. (2006); in fact, the FCC auctions are typically such iterative auctions (Cramton 1997, Federal Communications Commission 2006).

When iterative auctions are used to reduce the number of bundles considered in each round, the auctioneer still must solve a set partitioning problem in each round in order to determine the optimal allocation. Among others, Rothkopf et al. (1998), Bichler and

Kalagnanam (2005), Günlük et al. (2005), Sandholm et al. (2005), have presented heuristics and algorithms to solve these set-partitioning problems efficiently in specific applications. Müller (2006) and de Vries and Vohra (2003) explore how the structure of the bid matrix may in some cases result in winner determination problems which are very tractable set-partitioning problems.

Virtually all of this research is premised on the assumption that bidders are exogenously endowed with the valuation of each bundle – that is, there is little if any attention paid to the question of how bidders compute their bids. It is simply assumed that the values of the bids exist. The motivation for our research stems from the fact that there is often a clearly-defined structure underlying this determination. For example, a trucking company can determine their cost for carrying a bundle of loads by constructing the least-cost set of continuous moves that cover all of these loads (this can be found by solving a straightforward minimum cost flow problem). This cost is then the value that the trucking company would bid in a VCG auction (in other types of auctions, this would be adjusted – for example, in a first-price auction a markup would be added). Thus, our research focuses on understanding and exploiting this underlying cost-generating function which bidders use to determine their lowest cost associated with winning a bundle of items.

The idea of transmitting more information which helps the auctioneer figure out the price of bundles for the bidders is not new. Sandholm (2002) and de Vries and Vohra (2003) suggest an "oracle" model, where the auctioneer has access to an oracle which can reveal or compute bidders' costs on demand. More recently, Sandholm et al. (2006) demonstrate the real-life use of "expressive bidding", wherein bidders are allowed to provide more information which actually enables the auctioneer to construct bids by changing some attributes of the bid. The only work we are aware of that directly transmits the cost-structure in place of the bid is by Hobbs et al. (2000) in the context of energy auctions, although they do not directly address how the winner determination problem is solved. Boutilier (2002) uses a bidding language comprised of logical clauses, and shows empirically that the winner determination problem can be solved much faster than in a canonical combinatorial auction if the bidders' costs are amenable to being expressed using a small set of such clauses. Our work takes this line of research to its logical conclusion: the bidder should submit their entire cost structure, so that the auctioneer can implicitly construct bids when solving the winner determination problem.

We note that this idea cannot extend to all possible combinatorial auctions. The landmark work of Nisan and Segal (2006) shows that for general combinatorial auctions, in the worst case, an exponential number of bids must be transmitted if the optimal allocation is to be computed exactly (in the simplest example, if valuations were random). The difference is that in several real-world examples (such as the trucking example), the cost of a bundle is not arbitrary, and in fact stems from an underlying cost-generating function. Therefore the communication bound of Nisan and Segal (2006) does not apply to our work; conversely, our approach cannot be extended to any general combinatorial auction. Indeed, one of the main avenues of research opened by our work is the identification and formulation of special cases where our approach does work.

2.1 VCG Auctions

In this paper, we focus primarily on VCG mechanisms (Vickrey 1961, Clarke 1971, Groves 1973). We briefly describe them here, although the reader familiar with VCG auctions may skip to §3. For the single-item case, this mechanism is the familiar sealed-bid second price auction, wherein the buyer procures the object from the lowest bidder but pays the second-lowest bid. Such auctions are elegant and commonly studied because it is a dominant strategy for bidders to reveal their true cost – overbidding does not increase the winner's payment but decreases the likelihood of being awarded the item, whereas underbidding can lead to being awarded the item with a payment less than the cost to provide it. This truthful-bidding property eliminates the need for modeling strategic, competitive behaviors and greatly simplifies analysis.

The single-item VCG mechanism extends in a fairly straightforward manner to a combinatorial auction. In a combinatorial VCG auction, the payments are determined as follows. Suppose after the bidders have submitted their bids, we determine the optimal allocation, which costs z^*. This cost is obtained by summing over all bidders the costs of all bundles allocated to them. In this sum, suppose z_i^* is the cost incurred by bidder i. Bidder i therefore must be paid at least z_i^*, but also ought to be paid a premium (since if they always are only paid their cost, they might as well not participate in the auction). This premium is calculated as follows. All of bidder i's bids are removed from the data, but all bids by the other bidders are retained. The new optimal solution is computed, and let z_{-i}^* denote its cost. Clearly $z_{-i}^* \geq z^*$; this increase $(z_{-i}^* - z^*)$ is the premium awarded to bidder i. That is, bidder i is paid a total of $z_i^* + (z_{-i}^* - z^*)$. In effect, bidder i is paid their cost, along with a premium equal to the additional cost that the auctioneer would have incurred had bidder i been absent.

VCG auctions are well-studied for combinatorial auctions because they are incentive compatible and individually rational, i.e., they induce participants to reveal their true costs or valuations for bundles (Krishna (2002), ch. 16). Furthermore, the VCG auction is economically efficient (maximizes overall surplus), and among all efficient auctions, it minimizes the cost to the auctioneer (Krishna and Perry 1998) (under certain assumptions). The interested reader is referred to Ausubel and Milgrom (2006) and Krishna (2002) for further information about VCG auctions, their properties, and applicability for combinatorial auctions. VCG auctions are also recognized to have practical limitations (Ausubel and Milgrom 2006); in §4.4 we discuss privacy issues and in §4.5 we explore how bid-generating functions can be used in first-price or other combinatorial auction mechanisms.

3 Bid-generating Combinatorial Auctions for Truckload Procurement

In this section, we consider the problem of truckload procurement auctions as a way to demonstrate the details of how our methodology could be applied to a real-world problem. We initially present a simplified version of the truckload procurement auction, describing its bid-generating function and formulating its winner determination problem. We contrast this with the traditional, canonical implementation of the same problem, and demonstrate the benefits to be gained by using the bid-generating function directly, instead of actually enumerating the bids. We then outline enhancements to our initial formulation that would enable the capturing of more complex operational considerations. Finally, we identify the remaining questions and future research efforts that need to be addressed before the proposed approach could actually be implemented in practice.

Freight transportation plays a critical role in the U.S. economy. The Bureau of Economic Analysis (2005) reported over $300 billion in U.S. transportation expenditures in 2004. More than eighty percent of the U.S. transportation costs are for trucking, with over half of this made up of third party truckload carriers (Sheffi 2004). Given the importance of truckload transportation, truckload procurement auctions have attracted substantial research interest in recent years. Caplice and Sheffi (2006) provide a comprehensive exposition of the current state-of-the-art for procurement auctions for trucking services. Other such studies include those by Ledyard et al. (2002), Song and Regan (2003), Sheffi (2004), Sandholm et al. (2006) and Figliozzi et al. (2003).

Truckload carriers move full loads (i.e. trailers) directly from origin to destination, sometimes according to long-term contracts and other times as contracted through a spot market. The auction marketplace for single loads appears to be quite robust in practice (Internet Truckstop 2006, Huff 2006), bringing improved efficiencies to both truckers and shippers. When a single shipper has multiple loads to be covered, auctioning them simultaneously provides added benefits because truckers can achieve economies of scope by combining loads into tours that reduce empty mileage and therefore total cost. However, such combinatorial auctions have had relatively limited success; although a number of factors play into this, one of the most critical factors is the combinatorial burden associated with taking into account these interactions between loads (Caplice and Sheffi 2006, Ledyard et al. 2002, Sheffi 2004). For example, with only 20 loads, there are over a million combinations of loads for bidders to consider in order to fully realize the economies of scope; real world auctions can contain thousands of loads (Plummer 2003), resulting in a virtually infinite number of combinations.

Table 1: Cost characteristics of 3 truckers.

Carrier $i \in \mathcal{N}$	**Bid load cost for each load** $m^i_j,\ j = 1, 2, 3, 4$; (\$/mile)	**Estimated cost for each lane** $l \in \mathcal{L}$ e^i_{jk}, defined as:
1	110.00, 124.45, 110.00, 90.71; (1.1)	\$0.8 per mile
2	130.00, 147.08, 130.00, 107.20; (1.3)	\$1.0 per mile, except $e^2_{11} = e^2_{12} = 59$, $e^2_{21} = e^2_{22} = e^2_{41} = e^2_{42} = 98,\ e^2_{24} = e^2_{44} = 40$
3	120.00, 135.76, 120.00, 98.95; (1.2)	\$0.9 per mile, except $e^3_{23} = e^3_{43} = 12$

In practice, it is often the case that combinatorial auctions are run where only some subsets of loads are bid upon. Such auctions, however, are not guaranteed to capture the full economies of scope. The goal of our research (motivated by our experience with a large consumer goods manufacturer) is to address this challenge, and provide a methodology whereby the full economic value of bundling loads is realized, by both truckers and shippers.

3.1 Problem Definition

We begin with a simplified model of truckload procurement auctions. A single shipper conducts a procurement auction for a set of loads $\mathcal{M}$ that they require to be transported, with each load specified by an origin and a destination. The carriers (trucking companies) are the individual bidders in this auction, with the set of carriers denoted $\mathcal{N}$ and indexed by i. Each carrier has two sets of information governing its cost structure. First, for every load $j \in \mathcal{M}$, carrier i knows the direct cost m^i_j of moving load j from its origin to its destination, capturing for example, fuel, driver wages, truck depreciation, tolls, etc. Second, for every lane $l \in \mathcal{L}$ (here, $\mathcal{L}$ is the set of all lanes: i.e., movements from one location to another in the network), each carrier has an *estimated* cost e^i_l of repositioning the truck across lane l. On one extreme, e^i_l may be zero, reflecting an existing contracted load along lane l. At the other extreme, e^i_l may reflect the need to move empty across lane l. In general, e^i_l may be anywhere between these two extremes, capturing the carrier's estimated potential for obtaining loads along lane j, through other contracts or on the spot market (see Caplice (1996) for cost models involving uncertain future loads, for example, based on random arrivals, dwell times, etc.). Since lanes $l \in \mathcal{L}$ are only used for repositioning trucks from the destination of one load to the origin of the other, we often represent the lane l by the pair (j, k) of loads with load j's destination at the origin of lane l and load k's origin at the destination of lane l. Likewise, often the e^i_l is represented as e^i_{jk} to reflect this fact.

As a running example which we use to illustrate current auction practice as well as our

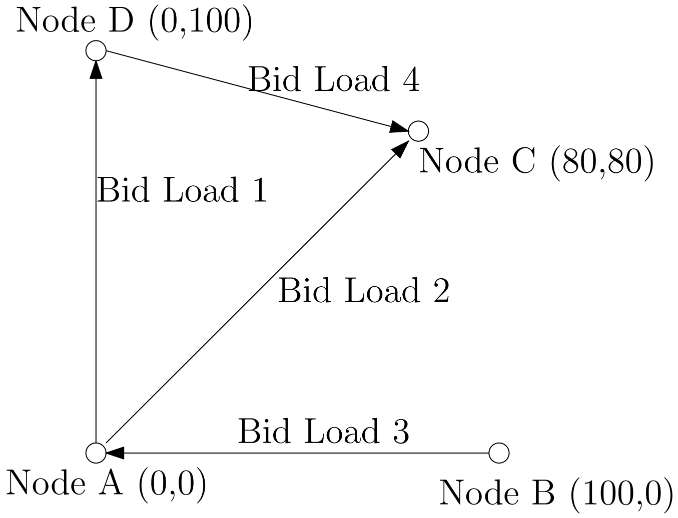

Figure 1: Initial network for truckload services procurement (x, y co-ordinates of the cities in parentheses).

proposed approach, consider the network of four cities in Figure 1. There are four loads up for auction, indicated by the solid lines; this constitutes the set $\mathcal{M}$. The cost structure of each carrier is shown in Table 1. For the bid loads, each carrier's cost m_j^i is just their loaded per-mile cost (1.1, 1.3 and 1.2 for the 3 carriers respectively) multiplied by the Euclidean distance. For the backhauls (i.e. the links from the destination of one load to the origin of another), the costs e_{jk}^i are for the most part given by the empty per-mile cost (0.8, 1.0, and 0.9 for the three carriers respectively) multiplied by the distance. However, there are some exceptions noted in the table. For instance, carrier 2 has a cost of only \$59.00 (instead of \$0.80$\times$ 100=\$80.00) from node D to node A, which could capture carrier 2's estimate that they have a fairly high chance of getting a future follow-on load along this lane on the spot market. This is reflected in e_{11}^2 and e_{12}^2, since node D is the destination of load 1 and node A is the origin of loads 1 and 2.

3.2 Canonical Combinatorial Auction (CCA) for Truckload Procurement Auctions

In the literature (Caplice and Sheffi 2006, de Vries and Vohra 2003), this procurement auction is run as follows. The shipper announces the set of loads $\mathcal{M}$. Each bidder i then submits a set of bids $S^i = \{(s_1^i, p_1^i), (s_2^i, p_2^i), \ldots, (s_{|S^i|}^i, p_{|S^i|}^i)\}$. Here $s_k^i \subseteq \mathcal{M}$ (i.e., one or more loads), and $p_k^i \in R$ is the price bid by bidder i to service the bundle of loads s_k^i. Note that the bidder may not (and in practice almost never) submits bids for all possible bundles of loads;

Table 2: Bids provided by 3 truckers.

Carrier	**Bundle of loads**	**Bid**
1	{1}	190.00
	{2}	214.96
	{3}	190.00
	{4}	156.68
	{2,3}	300.42
2	{1}	189.00
	{2}	245.08
	{3}	230.00
	{4}	147.20
	{1,3,4}	465.20
	{1,2,3,4}	694.74
3	{1}	210.00
	{2}	237.59
	{3}	210.00
	{4}	173.17
	{2,3}	267.76

that is, $|S^i|$ may be strictly less than $2^{|\mathcal{M}|} - 1$. The auctioneer (shipper) then solves the following set-partitioning problem, where variable x^i_k is set to 1 if bidder i is awarded the bundle of loads s^i_k and 0 otherwise:

$$\begin{aligned} z^*(\mathcal{N}) = \min \sum_{i \in \mathcal{N}} \sum_{k=1}^{|S^i|} p^i_k x^i_k & & & \\ \text{subject to} : \sum_{i \in \mathcal{N}} \sum_{k : m \in s^i_k} x^i_k & = & 1 & \qquad \forall m \in \mathcal{M} \\ x^k_i & \in & \{0, 1\} & \qquad \forall i \in \mathcal{N},\ k \in \{1, 2, \ldots, |S^i|\}. \end{aligned}$$

Each bidder i is then awarded the bundles for which $x^i_k = 1$. Since we are considering VCG auctions, p^i_k is the true cost to serve bundle s^i_k, and the actual payment to each bidder is given as follows. Let $z^*_i = \sum_{k=1}^{|S^i|} p^i_k x^i_k$ be the total cost of bids won by bidder i. Then bidder i is paid $z^*_i + (z^*(\mathcal{N} \setminus \{i\}) - z^*(\mathcal{N}))$.

In order to achieve the full benefits of CCA, each bidder must provide a bid for each combination of loads. In the example shown in Figure 1, each of the 3 carriers must provide 15 bids, since the 4 bid-loads can be used to create $2^4 - 1 = 15$ non-empty combinations.

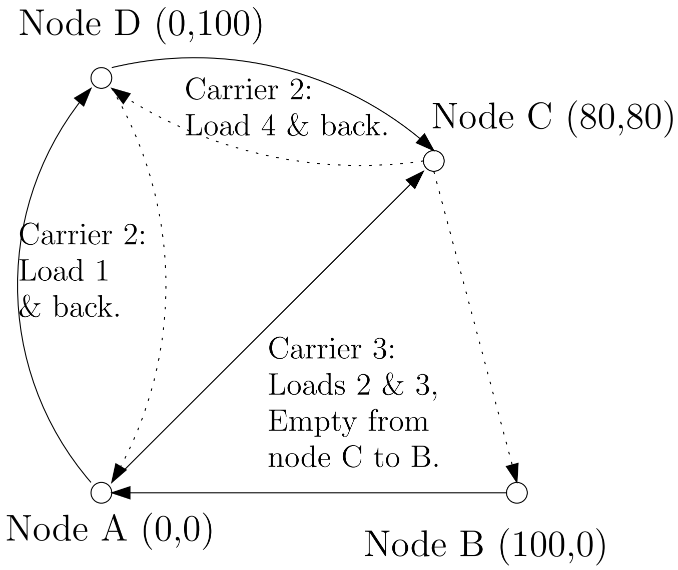

Figure 2: Allocation determined by CCA.

In practice, in all but the smallest auctions, it is impossible to enumerate all $2^{|\mathcal{M}|} - 1$ bids. Instead, carriers submit only a partial list of loads and bundles that they are interested in. In fact, Plummer (2003) reported that in a survey of 644 carriers, 72% of carriers only bid on single lanes, and even among those who bid on bundles, most submitted between 2 and 7 bundle bids. As a result, the solution quality is compromised.

In our example, suppose the bids received were as in Table 2. As the example shows, the bidders selectively provided bids for a small number of bundles. To determine the optimal allocation (with respect to the set of bids submitted), we now solve $(E-WDP)$ using these bids. The optimal solution is shown in Figure 2. Carrier 2 is awarded loads 1 and 4, with backhauls incurred by the carrier for each of those loads. Carrier 3's combination bid for the bundle $\{2, 3\}$ is also accepted, along with the associated backhaul from node C to node B. The total cost is 603.97. The actual payments to carriers 2 and 3 are 336.20 and 267.76 respectively, resulting in a total cost to the auctioneer of 647.08. As we shall see in §3.4, this is not the lowest cost solution, due to the fact that bidders did not submit bids on all possible bundles. The lowest cost solution actually costs 621.20, which would have been found had all bidders bid on all bundles.

Some attempts have been made to improve the quality of combinatorial auctions when not all bids can be enumerated. For example, bidders may specify minimum and/or maximum volumes for either specific lanes or their entire contract. They may also use logical combinations of bids; for example, specifying that they can serve "either load 1, or load 2 but not both", or "either load 3 or the combination of loads 4, 5 and 6". Caplice and Sheffi (2006) provide an excellent survey of combinatorial trucking auctions in practice. While the bidding strategies discussed above (volume conditions, logical bids) allow bidders to express some more information beyond listing a few bundles, they are not comprehensive either. Therefore, the set of bids received by the auctioneer is still incomplete. Furthermore, even if a fully enumerative set of bids could be communicated, the resulting exponentially large set partitioning problem would be intractable.

3.3 Cost Computation by Carriers in Canonical Combinatorial Auction

The defining question that motivates our approach is the following: *For a given carrier i and bundle s_k^i of loads, how does the carrier compute the cost of serving that bundle, p_k^i?* In the related auctions literature (Krishna 2002, Cramton et al. 2006), these valuations are considered "exogenously endowed"; that is, these numbers are available to the truckers from some external source, and the question of how the truckers come up with these numbers is not addressed.

Instead, we delve into the source of these costs p_k^i. In practice, if a trucker is offered a set of loads s_k^i and asked to quote a price for it, the trucker will evaluate how best that set of loads can be served, given the trucker's current cost and network structure. If the trucker's current cost and network structure is captured by the numbers $\{m_j^i\}_{j\in\mathcal{M}}$ and $\{e_l^i\}_{l\in L}$, the trucker computes p_k^i by solving the following optimization problem. Let x_j^i be a parameter that has value 1 if load j belongs to the given set of loads s_k^i and 0 otherwise. (This notation ties-in with our formulation in §3.4.) Define variables y_{jh}^i for $j, h \in s_k^i$, which takes value 1 if the carrier repositions from the destination of load $j \in s_k^i$ to the origin of load $h \in s_k^i$ and 0 otherwise. The objective function (1) sums the cost of the loaded moves and the connections between these moves. Constraints (2) and (3) state that for each bid load in the set, connection arcs to and from this load must be chosen.

$$p_k^i = \sum_{j \in s_k^i} m_j^i x_j^i + \min \sum_{j \in s_k^i} \sum_{h \in s_k^i} e_{jh}^i y_{jh}^i \tag{1}$$

$$\text{subject to} \qquad x_j^i - \sum_{h \in s_k^i} y_{hj}^i = 0 \qquad \forall j \in s_k^i \tag{2}$$

$$x_j^i - \sum_{h \in s_k^i} y_{jh}^i = 0 \qquad \forall j \in s_k^i \tag{3}$$

$$y_{jh}^i \in \{0,1\} \qquad \forall j,h \in s_k^i. \tag{4}$$

For example, consider once again the network displayed in Figure 1. Suppose bidder 3 wanted to construct a bid for the bundle of loads $\{2,3\}$. The bidder would set up an optimization problem as defined above, with appropriate values for m^3 and e^3, and the x^3 vector set to $(0,1,1,0)$. The resulting solution sets y_{23}^3 and y_{32}^3 to 1 and all other y variables to zero, indicating that the bidder should create a tour that covers load 3, followed immediately by load 2, followed by a repositioning move from node C (the destination of load 2) to B (the origin of load 3). This has a cost of 267.76, which would be bidder 3's bid for the bundle of loads $\{2,3\}$, which is in fact the case in Table 2. Bidder 1's bid for the same bundle of loads, however, would be 300.42; this is higher despite bidder 1's lower per-mile costs, since bidder 1 does not have a cheap backhaul opportunity on the lane from node C to node B.

If the auction is conducted as defined in §3.2, each carrier would have to solve the above problem to compute bid prices p_k^i for each bundle of loads they are interested in. While there are only 15 bundles in this example, the number of bundles grows exponentially with the number of loads. Since real-life auctions often involve hundreds of lanes and thousands of loads, constructing bids for all bundles is ruled out. In practice, carriers do such computations for only a small number of bundles, and follow other strategies to convey bid information (Plummer 2003), causing the inefficiencies to which we have alluded.

3.4 Bid-generating Combinatorial Auction (BCA) for Truckload Procurement Auctions

The crux of our formulation is the following: instead of bidding pairs of bundles of loads and prices, the bidder submits their cost function (m^i, e^i) to the auctioneer, where we use the shorthands m^i and e^i to represent the vectors $\{m_j^i\}_{j \in \mathcal{M}}$ and $\{e_l^i\}_{l \in \mathcal{L}}$ respectively. Observe that both these vectors are sub-exponential in size (the first size is the number of loads available for bid, the second size is the number of load-pairs). The auctioneer then uses this information to construct the optimal allocation and determine the winners of the auction, by solving the following problem. (This eliminates the need for the carrier to compute p^i_k for each possible bundle, but instead rolls in the computation of the bundle costs implicitly into the auction.)

Our formulation has two sets of binary variables, extending the idea behind the formulation (1)-(4). The parameter x^i_j defined in §3.3 now becomes a variable which determines whether load j is awarded to bidder i ($x^i_j = 1$) or not ($x^i_j = 0$). The variable y^i_{jk}, as before, takes the value 1 if carrier i repositions from the destination of load j to the origin of load k and 0 otherwise, and is defined for $j, k \in \mathcal{M}$ (so that the pair $(j, k) \in \mathcal{L}$). The integer programming formulation now follows.

$$z^*(\mathcal{N}) = \min \sum_{i\in\mathcal{N}} \sum_{j\in\mathcal{M}} m^i_j x^i_j \quad + \quad \sum_{i\in\mathcal{N}} \sum_{j\in\mathcal{M}} \sum_{k\in\mathcal{M}} e^i_{jk} y^i_{jk} \tag{5}$$

$$\text{subject to} \qquad x^i_j - \sum_{k\in\mathcal{M}} y^i_{kj} \;=\; 0 \qquad \forall i \in \mathcal{N}, j \in \mathcal{M} \tag{6}$$

$$x^i_j - \sum_{k\in\mathcal{M}} y^i_{jk} \;=\; 0 \qquad \forall i \in \mathcal{N}, j \in \mathcal{M} \tag{7}$$

$$\sum_{i\in\mathcal{N}} x^i_j \;=\; 1 \qquad \forall j \in \mathcal{M} \tag{8}$$

$$x^i_j \;\in\; \{0, 1\} \qquad \forall i \in \mathcal{N}, j \in \mathcal{M} \tag{9}$$

$$y^i_{jk} \;\in\; \{0, 1\} \qquad \forall i \in \mathcal{N}, j, k \in \mathcal{M}. \tag{10}$$

In the formulation above, the constraints (8) ensure that each load up for bid is awarded to a carrier. The constraints (6) and (7) ensure "conservation": that is, if a truck is covering load j, it must move from the destination of that load to the origin of some other bid load, and conversely. Notice that if the origin of the subsequent load is the same as the destination of the first load j, then this repositioning move (of cost 0) is effectively non-existent and is purely a modeling aid. These constraints together ensure that all bid loads in $\mathcal{M}$ are covered as parts of "continuous moves" or "loops" by carriers, so that the truck returns to its starting point. The objective function (5) accounts for the cost of all the loaded moves up for bid, as well as the cost of all lanes that have to be traveled to complete round trips for the bid loads. Indeed, the purpose of this formulation is to figure out how truckers can best leverage their current outlook on existing and potential contracts while covering all the loads that are up for bid.

This formulation is in fact a special case of the well-known multicommodity flow problem

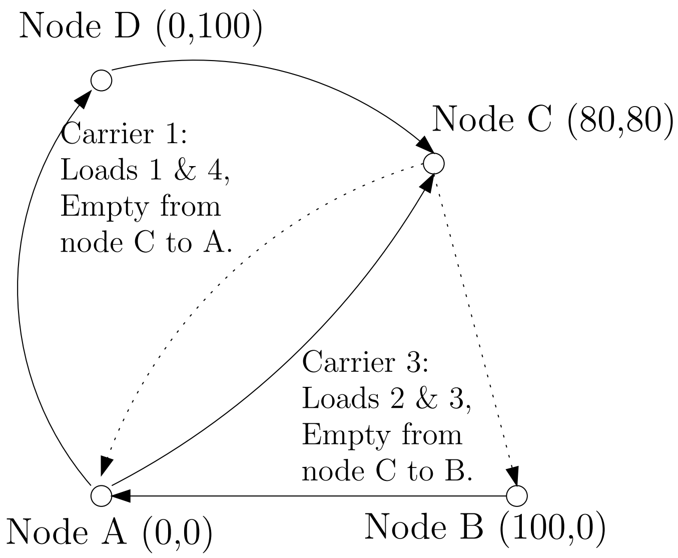

Figure 3: Optimal allocation determined by BCA.

which is typically easy to solve in practice even for large instances (Ahuja et al. 1993). Furthermore, the integrality of the y variables can be relaxed. This makes the problem even easier to solve. We state the relaxation claim below, and defer the proof to the appendix.

Proposition 1 *For any problem of the form (5)-(10), replacing the integrality constraint (10) with the non-negativity constraint* $y^i_{jk} \geq 0 \ \ \forall i \in \mathcal{N}, j, k \in \mathcal{M}$ *results in a problem with the same optimal solution as the original problem.*

As in §3.2, since we are using a VCG auction, the payments are computed as follows. The cost incurred by bidder i is $z^*_i = \sum_{j\in\mathcal{M}} m^i_j x^i_j + \sum_{j\in\mathcal{M}} \sum_{k\in\mathcal{M}} e^i_{jk} y^i_{jk}$. Then, the payment to bidder i is $z^*_i + (z^*(\mathcal{N} \setminus \{i\}) - z^*(\mathcal{N}))$.

The auction above is referred to as the **Bid-generating Combinatorial Auction**, abbreviated BCA, for the rest of this paper. The auction will be generalized for other problems in §4, but we will continue to use the term BCA to refer to it.

3.5 Illustration of BCA for Truckload Procurement Auctions

We continue with our example initiated in Figure 1, and show how BCA would solve the same problem. BCA begins by computing the optimal allocation by solving the program (5-10). This allocation is described below, and is shown in Figure 3.

Carrier 1 is awarded loads 1 and 4, incurring a total cost of 291.22, including the repositioning move from node C to node A. Carrier 3 is awarded loads 2 and 3, for a total cost of

Table 3: Comparison of CCA and BCA for trucking services procurement.

	CCA formulation	BCA formulation
No. of Integer variables	$N(2^M - 1)$	NM
No. of Continuous variables	0	NM^2
No. of Constraints	M	$M(2N+1)$

267.76 (as in CCA), including a backhaul from node C to node B, The VCG auction results in a total payment of 320.78 to Carrier 1 and 300.42 to Carrier 3, so that the auctioneer's total cost is 621.20.

Notice that the auctioneer's total cost under BCA (621.20) is lower than that under CCA (647.08) for our example. The reason is simple: In CCA, Bidder 1 did not provide a bid for the bundle of loads $\{1, 4\}$. Resultantly, those loads ended up being awarded to bidder 2, who has a higher cost structure. If all bidders had provided bids for all combinations (or in this case, if at least Bidder 1 had bid for the bundle $\{1, 4\}$), CCA would indeed have found the optimal allocation.

This illustrates the main benefit of our work – the ability to implicitly determine optimal bundles and allocations, without bidders having to explicitly bid on each of an exponential number of possible bundles, nor the auctioneer having to solve an exponentially large set partitioning problem. All of the existing methodologies in the truckload auction literature either require (implicit or explicit) potentially exhaustive enumeration of all possible combinations, or run the risk that the optimal allocation is not discovered due to some bundle not being bid upon (as was the case in our example of CCA).

3.6 Implications for Larger Scale Problems

Although having an underlying multicommodity flow structure in BCA is certainly a benefit, the key advantage of our approach stems from the reduction in size of the winner determination problem. Table 3 compares the size of our BCA formulation for this problem with the corresponding CCA formulation, where N is the number of bidders, and M is the number of loads. Observe that for all but the smallest number of loads, the difference in the number of integer variables is quite dramatic. For example, with 100 loads and 10 bidders, CCA would have approximately 10^{31} integer variables, while our formulation has a mere 1000.

3.7 Additional Constraints and Requirements

While the model shown in (5)-(10) suffices to cover all loads offered by the shipper, carriers may have other constraints that they would like to consider when they participate in such auctions. We briefly explore below how some such constraints can be modeled in our formulation. The actual integration of these constraints into the auction methodology is a separate, ongoing, research project.

- **Carrier capacity**: Suppose carrier i had a restriction of a maximum of μ_i on the total miles that they can travel (due to fleet size, federal regulations, and other such considerations). This can be modeled as follows, where c_j for a load denotes the mileage for that load, and c_{jk} is the distance in miles from the end of load j to the origin of load k:

$$\sum_{j \in \mathcal{M}} c_j x_j^i + \sum_{j \in \mathcal{M}} \sum_{k \in \mathcal{M}} c_{jk} y_{jk}^i \leq \mu_i.$$

  Carrier i may also have a restriction that if they are to serve this shipper, it must be for a certain pre-specified minimum number of miles. Such a constraint can be modeled in a similar fashion.

  Likewise, if the shipper wants to impose minimum and/or maximum limits on the amount of miles allocated to any single carrier, this can be done using appropriate modifications to this constraint. Shippers sometimes desire such constraints in an attempt to provide sufficient business to each carrier so as to make it a meaningful relationship and achieve economies of scale (necessitating a minimum) while also alleviating risk due to a single carrier being responsible for too much of the shipping (necessitating a maximum).

- **Number of winners**: The shipper might want to award loads to no fewer than $i_{\min}$ and no more than $i_{\max}$ carriers, in order to make the relationships manageable while spreading the risk. Such a restriction can be modeled as follows, where we introduce a new variable w_i which takes the value 1 if carrier i is awarded a load and zero otherwise:

$$x_j^i \leq w_i, \quad \forall j \in \mathcal{M},\ i \in \mathcal{N}$$

$$i_{\min} \leq \sum_{i \in \mathcal{N}} w_i \leq i_{\max}.$$

- **Round-trip mileage limits**: One of the biggest concerns of truck drivers is the frequency with which they are able to return home. For the trucking industry, this translates to constructing short round-trips, ensuring that not too many loads are strung together to form a very long continuous move round-trip.

  Modeling and implementing this requirement is a challenging research problem in itself, and is currently on-going work. However, we note that our formulation explicitly outputs tours. Every load j, when assigned to a carrier i by setting x_j^i to one, also results in one connecting load preceding it ($y_{kj}^i = 1$ for some k), and one connecting load succeeding it ($y_{jk}^i = 1$ for some k). These variables can be used to uniquely construct tours for every load and every carrier. Once these tours are constructed, one can pursue several methods of enforcing the tour-length constraint (such as branch-and-bound approaches), which is the subject of future research.

4 Generalization

The key idea in the truckload procurement auction example is that the cost of any set of loads can be constructed by building a least-cost set of tours to cover these loads. The tours can easily be found by solving a minimum cost flow problem using the per-mile cost associated with the bid loads plus the estimated cost of moving between any pair of bid loads. Thus, it is sufficient for the bidder to provide a set of arc costs, rather than enumerating all bundles of loads, solving a minimum cost flow problem for each bundle, and then providing this exponentially large set of bids to the auctioneer. Similarly, the auctioneer can explicitly embed each bidder's arc costs in a master multicommodity flow problem (i.e. one minimum cost flow problem per bidder, linked to each other by the need to cover each bid load exactly once), rather than solving an exponentially large set partitioning problem. The results of this multicommodity flow problem in turn reveal the assignment of loads to bidders.

In this section, we present a generalization of this approach and discuss its properties, laying the groundwork for other applications discussed in §5, as well as the research questions discussed in §6. We continue to focus on multi-item, single round, procurement auctions. However, our results also apply to forward auctions where objects are being sold; in particular, the wireless spectrum auction we discuss in §5.1 is such an auction.

The auction is conducted by a single buyer, with a set $\mathcal{N} = \{1, 2, \ldots, N\}$ of prospective sellers (bidders). The set $\mathcal{M} = \{1, 2, \ldots, M\}$ is the set of objects/services being auctioned. A bundle or combination is a subset of $\mathcal{M}$, and different bidders are interested in selling different such bundles at different prices. There are $2^M - 1$ such bundles. The auctioneer's objective is to procure one item of each of the objects in $\mathcal{M}$ at the lowest possible total price.

Let bidder i's cost for bundle S be given by $c_i(S)$. (In the trucking example, $c_i(S)$ was called p_k^i, where the bundle S corresponds to the set of loads s_k^i.) As shown in §3.3 for the trucking example, the bidder's cost for this bundle is often the solution to some optimization problem. The key idea of our approach is that instead of the bidder having to solve this optimization problem for a large set of bundles, the bidder just transmits the parameters of their optimization problem to the auctioneer as their "bid". The auctioneer then simultaneously, in a single round, solves one large optimization problem which determines the allocation of goods to bidders as well as the true costs to each bidder of the bundle awarded to them.

This approach was illustrated in the previous section in the context of truckload procurement auctions, and is generalized in the second half of this section. First, we establish the basics of a general combinatorial auction.

4.1 Canonical Combinatorial Auction

VCG combinatorial auctions are canonically conducted in a manner which is a straightforward generalization of the process discussed in §3.2, along with the VCG framework discussed in §2.1. Recall that bidder i has a specific cost function c_i, and the cost to bidder i of providing bundle S is given by $c_i(S)$. In its most elementary form, the auction proceeds as follows:

CCA **(Canonical Combinatorial Auction)**

1. Each bidder provides costs $c_i(S)$ for every bundle $S \in 2^{\mathcal{M}}$. We require the cost function to be non-negative (a natural assumption). Initially, we also require the cost function to be sub-additive; that is, for every $S \subset \mathcal{M}$ and $j \in \mathcal{M} \setminus S$, and for every bidder i, we have $c_i(S \cup \{j\}) \leq c_i(S) + c_i(\{j\})$. Sub-additivity is a natural assumption stating that the union of two bundles costs no more than the sum of the two bundles taken separately. We remark on the implications of sub-additivity later in this section. We also observe that the cost function in the trucking example can be shown to be sub-additive (Proposition 4 in the Appendix).

2. The auctioneer solves a winner-determination problem, where she partitions the objects in $\mathcal{M}$ and awards bundles to bidders in such a way that the total cost is minimized. That is, the auctioneer solves the following mathematical program, denoted $E-WDP$, where the E stands for "enumeration" and the binary variable $x_i(S)$ takes the value 1 to indicate that bidder i is awarded bundle S, and 0 otherwise.

$$z^*(\mathcal{N}) = \min \sum_{i \in \mathcal{N}} \sum_{S \subseteq \mathcal{M}} c_i(S) x_i(S) \qquad (E-WDP)$$

$$\begin{aligned} \text{subject to}: \sum_{i \in \mathcal{N}} \sum_{S \subseteq \mathcal{M}: j \in S} x_i(S) &= 1 && \forall j \in \mathcal{M} \\ x_i(S) &\in \{0,1\} && \forall i \in \mathcal{N},\ S \subseteq \mathcal{M}. \end{aligned}$$

   Sub-additivity ensures the existence of an optimal solution awarding no more than one bundle per bidder. For notational concision, we denote the cost incurred due to bidder i as z_i^*. That is, $z_i^* = \sum_{S \subseteq \mathcal{M}} c_i(S) x_i^*(S)$.

3. In order to determine the payment to bidder i, the auctioneer computes the surplus added to the system by bidder i. That is, the auctioneer begins by computing $z^*(\mathcal{N} \setminus i)$,

which is the total cost of providing all the bundles had bidder i been absent. The difference $z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N})$ is the premium awarded to bidder i, so that the payment he receives is $z_i^* + z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N})$.

For the rest of this paper, we use the abbreviation CCA to refer to the auction methodology described above. This traditional VCG combinatorial auction suffers due to the combinatorial explosion in the number of bundles (an exponential function of the size of the set $\mathcal{M}$ being auctioned), presenting a significant hurdle to its practical implementation.

Researchers have suggested several approaches to ameliorate this hurdle for specific contexts, as surveyed in §2. These suggested approaches either sacrifice capturing the entire cost function c_i (losing cost optimality), and/or require several auction rounds (resulting in increased procurement cycle times). In contrast, our proposed methodology provides the exact optimal allocation and payments in a single round, and more importantly avoids combinatorial explosion in the number of bids and winner determination problem size. We formally specify and compare our proposed methodology to CCA in the following section.

Note that even if non-VCG mechanisms are considered, the problems caused by the combinatorial explosion remain. Furthermore, if the non-VCG auction possesses well-defined bid-generating computations on the part of bidders, our methodology is still applicable, as discussed further in §4.5.

4.2 Proposed Methodology: Bid-generating Combinatorial Auction

Our proposed methodology eliminates the problems caused by the combinatorial explosion for those classes of problems where the cost functions of bidders for bundles can be embedded in tractable winner determination problems. As noted above and shown in the truckload procurement auction, we require that the cost function $c_i(S)$ of bidder i providing bundle S be computable as the optimal solution to some optimization problem. That is,

$$c_i(S) = \min_{y \in P} \tilde{c}_i^T(x_S, y) \text{ s.t. } A_i(x_S, y) \leq b_i. \tag{11}$$

In the equation above, the matrix A_i and vector b_i are the bidder-specific parameters which determine the bidder's cost structure. To place this in the context of the truckload example, A_i and b_i are generated from the constraints (2) and (3). The y variables are "linking" variables, which allow the bidder to model the cost of the bundle accurately; in the trucking example, the y variables played the role of indicating which lanes the trucks would have to travel to complete the round-trips for the bid loads, thereby affecting costs. There is also a cost vector $\tilde{c}_i$ which is bidder-specific; in the trucking example, this cost vector was the concatenation of the vectors m^i and e^i. Finally, the restriction $y \in P$ enforces additional feasibility requirements, such as the integrality constraints (4) in the trucking example. The vector x_S is simply an indicator vector, where the j^{th} component takes the value 1 if object $j \in \mathcal{M}$ belongs to bundle S and 0 otherwise; this plays the same role as the x vector in the formulation in §3.3.

There are three components in the equation above that are bidder-specific: $\tilde{c}_i, A_i$, and b_i. We combine these three into a single structure, θ_i, which contains sufficient information to construct $\tilde{c}_i, A_i$ and b_i. We refer to θ_i as the *true cost type* of bidder i. Our proposed auction is then defined as follows:

Bid-generating Combinatorial Auction (BCA)

1. Each bidder provides its true cost type θ_i.

2. The auctioneer solves the winner determination problem (now denoted $T-WDP$, for cost Type) defined as follows. Here the variable x_{ij} takes the value 1 if bidder i is awarded object $j \in \mathcal{M}$ and 0 otherwise.

$$z^*(\mathcal{N}) = \min \sum_{i \in \mathcal{N}} \tilde{c}_i^T(x_i, y_i) \qquad (T-WDP)$$

$$\begin{aligned} \text{subject to:} \sum_{i \in \mathcal{N}} x_{ij} &= 1 && \forall j \in \mathcal{M} \\ A_i(x_i, y_i) &\leq b_i && \forall i \in \mathcal{N} \\ x_{ij} &\in \{0,1\} && \forall i \in \mathcal{N},\ j \in \mathcal{M} \\ y_i &\in P_i && \forall i \in \mathcal{N}. \end{aligned}$$

   As before, we use $z_i^* = \tilde{c}_i^T(x_i^*, y_i^*)$ to denote the cost incurred by bidder i in the optimal solution.

3. The payment to bidder i, as before, is $z_i^* + z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N})$.

Compared to the canonical combinatorial auction, our approach has several significant benefits:

Benefit 1. The cognitive burden on the bidder is greatly reduced – rather than computing and communicating an exponential number of bids, only the true cost type must be relayed.

Benefit 2. The computational burden on the auctioneer is greatly reduced by eliminating the need to solve an exponentially large integer program, and instead solving a much smaller program. While the resulting problem is still an integer program and thus in the worst case can take exponential time, practical instances of such problems are routinely solved in acceptable timeframes.

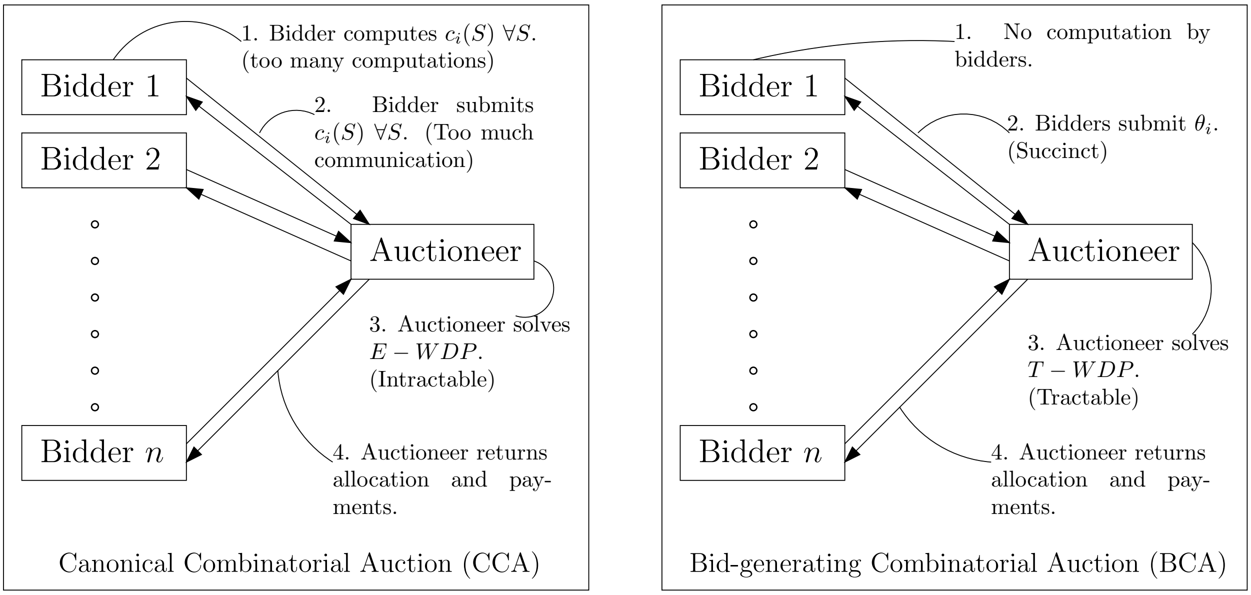

Figure 4: Comparison of CCA and BCA

Benefit 3. BCA is equivalent to the fully-enumerated VCG auction CCA (as proved below in Proposition 2) and yields all of the corresponding benefits.

Figure 4 highlights these benefits, contrasting the canonical combinatorial auction with BCA.

Proposition 2 *Consider a combinatorial VCG auction for the set of objects $\mathcal{M}$, where for each bidder $i \in \mathcal{N}$ the cost to provide the bundle $S \subseteq \mathcal{M}$ is given by $c_i(S)$. Suppose there exists a type θ_i for each bidder and functions $\tilde{c}_i = \tilde{c}(\theta_i)$, $b_i = b(\theta_i)$ and a matrix $A_i = A(\theta_i)$, such that for every $S \subseteq \mathcal{M}$, we have $c_i(S) = \min_y \tilde{c}_i^{\,T}(x, y)$ s.t. $A_i(x, y) \leq b_i$, where x is the incidence vector for S; that is, x is a binary vector with M components, with the j^{th} component being 1 if $j \in S$ and 0 otherwise. Then, for any set of types $\{\theta_i\}_{i\in\mathcal{N}}$, the optimal solutions of the integer programs $(E-WDP)$ and $(T-WDP)$ have the same values.*

Proof: Let $x^*(E)$ be an optimal solution to the formulation $(E-WDP)$ of cost $z^*(E)$, and let $(x^*(T), y^*(T))$ be an optimal solution to $(T-WDP)$ of cost $z^*(T)$. Observe that $z^*(T) = \sum_{i\in\mathcal{N}} z_i^*(T)$, and likewise $z^*(E) = \sum_{i\in\mathcal{N}} z_i^*(E)$.

First, consider the solution $x^*(E)$. Let $S_i^*(E)$ be the set of goods allocated to bidder i in this solution. Therefore, $z^*(E) = \sum_{i\in\mathcal{N}} c_i(S_i^*(E))$. Define the vector $\hat{x}_i(T)$ to be the incidence vector of $S_i^*(E)$; that is, $\hat{x}_{ij}(T) = 1$ if $j \in S_i^*(E)$ and 0 otherwise. By definition, there must exist a vector $\hat{y}_i$ such that $c_i(S_i^*(E)) = \tilde{c}_i^{\,T}(\hat{x}_i(T), \hat{y}_i)$ and $A_i(\hat{x}_i(T), \hat{y}_i) \leq b_i$. Let the vector $\hat{x}(T)$ be defined as the concatenation of the $\hat{x}_i(T)$ vectors for each bidder, and similarly define $\hat{y}$ as the concatenation of the $\hat{y}_i$ vectors. The covering constraint in $(T-WDP)$ continues to hold, so that $(\hat{x}(T), \hat{y})$ is a feasible solution to $(T-WDP)$, of cost $z^*(E)$. Since the optimal solution to $(T-WDP)$ can only do better, we must have $z^*(T) \leq z^*(E)$.

Conversely, consider the solution $(x^*(T), y^*(T))$. For any bidder i, let $S_i^*(T)$ be the bundle allocated to bidder i in BCA, and define the vector $x(E)$ appropriately where $x_i(S) = 1$ for $S = S_i^*(T)$ and 0 otherwise. By definition, we have $c_i(S_i^*(T)) \leq \tilde{c}_i^T(x_i^*(T), y_i^*(T)) = z_i^*(T)$. Therefore, the solution given by $(x(E))$ is feasible for $(E-WDP)$, and has cost no more than $\sum_{i \in \mathcal{N}} z_i^*(T) = z^*(T)$. The optimal solution to $(E-WDP)$ can only do better, so $z^*(E) \leq z^*(T)$.

This proves that $z^*(T) = z^*(E)$, which is the claim of the proposition. $\square$

Corollary 1 *The premium paid to each bidder is the same under both* CCA *and* BCA *auctions.*

Proof: Recall that the premium paid to bidder i is $z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N})$. Proposition 2 guarantees that $z^*(\mathcal{N})$ is the same under both CCA and BCA, and so is $z^*(\mathcal{N} \setminus i)$. The corollary now follows. $\square$

Corollary 2 *If there is a unique optimal allocation in one of the two solution methodologies (*CCA *or* BCA*), then the same allocation is also the unique optimal allocation in the other auction, and each bidder gets the same total payment in both auctions.*

Proof: Observe that the proof of Proposition 2 relies on showing that the optimal allocation of one auction is a feasible allocation in the other auction of same or lower cost. Therefore, if the optimal allocation is unique in one auction, the same allocation must be the unique optimal allocation in the other. The cost incurred by bidder i in both auctions is therefore the same, z_i^*. Since the payment to each bidder is his cost plus his premium, Proposition 2 and Corollary 1 imply that the payment to each bidder is the same under both auctions. $\square$

Proposition 3 BCA *is an individually rational and incentive compatible auction.*

Proof:

The cost incurred in $(T-WDP)$ by agent i is z_i^*; by truthful bidding, z_i^* is also i's true cost for the bundle he is awarded. Hence the payment minus cost (profit) of agent i is simply $z_i^* + z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N}) - z_i^*$, or $z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N})$. The value $z^*(\mathcal{N} \setminus i)$ must be no less than $z^*(\mathcal{N})$ since the optimal allocation in $\mathcal{N} \setminus i$ is feasible and has the same cost even with bidder i included. Hence $z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N})$ is always non-negative, implying individual rationality. (As is convention, bidders' reservation profits are normalized to zero.)

We next prove incentive compatibility. Let $z^*(\mathcal{N})$ be the optimal value of $(T-WDP)$ over agent set $\mathcal{N}$ with truthful reporting of types, and let $\hat{S}_j$ be the optimal allocation to agent j specified by $(T-WDP)$ when agent j reports type $\hat{\theta}_j$, for $j \in \mathcal{N}$. Note that the cost incurred to the auctioneer by agent j in this optimal allocation is $z_j^* = c(\hat{S}_j, \hat{\theta}_j)$, and j's true cost for his bundle is $c(\hat{S}_j, \theta_j)$. Suppose that $\hat{\theta}_j = \theta_j$ for all j except possibly i. Then agent i's profit (payment minus cost) is $c(\hat{S}_i, \hat{\theta}_i) + z^*(\mathcal{N} \setminus i) - \sum_{j \in N \setminus i} c(\hat{S}_j, \theta_j) - c(\hat{S}_i, \hat{\theta}_i) - c(\hat{S}_i, \theta_i)$, which equals $z^*(\mathcal{N} \setminus i) - \sum_{j \in N} c(\hat{S}_j, \theta_i)$. The first term is independent of $\hat{\theta}_i$. The second term is the objective function of $(T-WDP)$ with truthful reporting evaluated at some feasible allocation $\hat{S}_j,\ j \in \mathcal{N}$, and therefore is no less than the optimal value $z^*(\mathcal{N})$. Hence, if all bidders $j \neq i$ report their true cost, i's profit is no greater than $z^*(\mathcal{N} \setminus i) - z^*(\mathcal{N})$. Since this profit is achieved when he reports his true cost θ_i, we are done.

□

Therefore, it is a dominant strategy for every bidder to reveal their type vector θ_i truthfully. We discuss other issues regarding the practical implementation of BCA in §4.4.

4.3 Sub-Additivity in BCA

Recall that we initially assumed that our cost function is sub-additive; that is, the cost of a combination of two bundles is no more than the sum of costs of the bundles taken individually. While this is often a natural assumption, it need not always hold. For example, if a bidder has a capacity constraint and can provide no more than 3 items, then the cost of a bundle of 4 items is effectively infinity, which is greater than the sum of costs of two bundles of 2 items each.

If the cost function is not sub-additive, in CCA, one would add the following constraint:

$$\sum_{S \subseteq \mathcal{M}} x_i(S) \leq 1 \quad \forall i \in \mathcal{N}.$$

Note, however, that this sums over an exponential number of variables.

In BCA, it may instead be possible to add application-specific constraints to the set $A_i(x_i, y_i) \leq b_i$ to enforce these requirements. For example, if sub-additivity is violated due to a capacity constraint of the form that the bidder can provide no more than 3 items, then we simply add the constraint $\sum_{j \in \mathcal{M}} x_{ij} \leq 3$ to $A_i(x_i, y_i) \leq b_i$. Indeed, the carrier capacity constraint in §3.7 is one such constraint, as are the capacity constraints in §5.3.

4.4 Privacy Issues of BCA

Since BCA is incentive compatible and individually rational (Proposition 3), it is in a bidder's best interests to bid truthfully. Nonetheless, bidders might be understandably skeptical about providing private information that could be of value to their competitors if divulged.

Likewise, from the auctioneer's standpoint, there is the risk of false-name bidding and other forms of collusion.

Such concerns are not insurmountable, however, and we suggest that our approach provides incentive to overcome these obstacles so as to achieve the economies of scope of a fully-enumerated combinatorial auction. Consider, as an analogous example, the related data disclosure issues of a Vendor Managed Inventory (VMI) system. Such systems require that companies supplying goods to retail stores be given direct visibility into the retailers' inventory levels of their goods (Chopra and Meindl (2007), pp 518). Like our proposed auction approach, VMI systems also raise concerns about data security and privacy. Largely due to the anticipated benefits of these systems, however, vendors such as Wal-Mart have overcome these concerns and VMI systems are now actively used in widespread practice (Lee et al. 1997).

In the case of auctions, trusted third-party companies such as CombineNet and Manhattan Associates, can (and do) provide services such as bidder pre-qualification and transaction confidentiality, helping to reduce the risk of collusion and the leaking of data. There is also an emerging stream of literature in cryptographically-secured auctions (Kudo 1998, Franklin and Reiter 1996) which focuses on data privacy as well. Although we do not wish to trivialize bidders' privacy concerns, we nonetheless suggest that if the benefits to bidders of releasing their true-cost types is sufficient, then methods to overcome their concerns will be found.

4.5 BCA in a First Price Mechanism

Throughout this manuscript, we have assumed a VCG mechanism. This is because the dominant bidding strategy in VCG mechanisms is for bidders to bid truthfully, and thus the bid-generating function is simply the underlying cost function that bidders use to value bundles. Our approach does not require truthful bidding (or a VCG-type mechanism), however; we simply use this mechanism for its clarity and ease of exposition.

In fact, our approach extends to any auction mechanism, so long as the bidders have a succinct bid-generating function. In particular, consider the first-price auction. In the canonical version of this auction, bidders submit bids for bundles which already include a profit markup. The auctioneer then solves the winner determination problem $(E - WDP)$, and pays bidder i an amount $z_i^* = \sum_{S \subseteq \mathcal{M}} c_i(S) x_i^*(S)$ as defined in §4.1. Bidders still have to compute their bids for these bundles, and there is some research exploring good bidding strategies for such combinatorial auctions (An et al. 2005, Günlük et al. 2005).

In the BCA version of a first-price combinatorial auction, all we require is that the bidding strategy be representable by succinct bid-generating functions (as is the case in An et al. (2005) and Günlük et al. (2005)). If θ_i represents the parameters of the first-price bid-generating function for bidder i, then in the first price version of BCA, each bidder simply transmits these parameters θ_i. The auctioneer then solves $(T - WDP)$ as defined in §4.2, with the payment to bidder i given by $z_i^* = \tilde{c}_i^T(x_i^*, y_i^*)$.

5 Other Application Areas

The success of implementing BCA in other applications depends on whether the bid-generating function can be used to formulate a tractable mathematical program of the winner determination problem. In this section, we show how such a mathematical program can be constructed for three other well-known applications of combinatorial auctions: wireless spectra, energy markets, and operations procurement. Observe that the energy market auction is actually a forward auction where items are sold, thus demonstrating that our approach is applicable in both forward and reverse auctions.

5.1 Wireless Spectrum Auctions

The auction of frequencies for wireless communication has been one of the most dramatic and insightful applications of combinatorial auctions. They have been held in many countries for different industry segments, and continue to be held: Most recently, in 2006, the United States Federal Communications Commission auctioned frequencies to allow in-flight cellphone usage in passenger aviation. Wireless spectrum auctions provide ample scope for combinatorial bids: Bidders gain from scale, as well as from contiguity of regions where they win presence and contiguity of frequencies within a single region (Günlük et al. 2005). Currently, the mechanism most favored by the FCC is an iterative combinatorial auction (Federal Communications Commission 2006), where bidding goes on for several rounds determined by a specific set of rules (in this particular case, 161 rounds over 29 days resulting in 104 winning bidders for 1087 licenses). BCA has the potential advantage of making the auction a single-round mechanism, while at the same time capturing all the economic efficiencies of a enumerative combinatorial auction.

For illustrative purposes, in this section we will focus on a simple auction of frequencies for wireless cellphone communications (others include in-flight communications, land-lines, broadcast spectra, etc.). In such an auction, the unit commodity is a frequency band in one specific location (such as a metropolitan area or state). Several locations are up for bid (typically encompassing the entire nation), and multiple frequency bands are also offered for auction in each location.

While the valuation by each bidder of a bundle of licenses could be a fairly complex function in general, the literature (Cramton 1997, Günlük et al. 2005) suggests that the two biggest combinatorial components of the valuation are synergies in locations allocated and frequencies allocated. Although other synergies can also be modeled in our framework, in order to keep the discussion straightforward and brief we choose to focus on locational synergies. In the spirit of the estimated costs in the trucking example of §3.1, we assume the bidders can compute the values of such synergies.

Locational synergies occur when the same bidder wins licenses for a set of locations which together give the bidder more value than the sum of the individual valuations. For example, if a bidder wins licenses for the metro areas containing New York and Philadelphia, they may benefit from becoming a major player in that region. Therefore the bidder may be willing to pay more for the bundle of these two locations than the sum of the bidder's valuations for winning each of these in isolation. Locational synergies may also occur if bidders already have licenses in some areas and only want to cover their gaps, or they may want to focus on large metropolitan areas, etc.

5.1.1 Mathematical Formulation

Note that although our initial model in §4 is in terms of a reverse auction, this application is a regular auction where bundles are sold. However, the notation and formulation carry over with minimal changes. Let $\mathcal{N}$ denote the set of bidders (wireless providers), indexed by i. Let $\mathcal{L}$ denote the set of locations, indexed by l. Thus the individual items being auctioned are the locations in $\mathcal{L}$. We now show one possible formulation as a multi-item combinatorial auction which allows bidders to leverage the locational synergies.

Since this is a forward auction, the locational synergies actually imply that super-additivity holds across the units being auctioned, which is precisely what we desire. That is, if $L' \subset \mathcal{L}$ represents a set of locations and $v(i, L')$ represents bidder i's valuation for the bundle L', then $v(i, L') \geq \sum_{l \in L'} v(i, l)$. In general, the valuation function v for any bidder i must now be specified over all subsets of $\mathcal{L}$, which is precisely the combinatorial explosion BCA enables us to avoid.

The key concept to avoid the combinatorial explosion is to associate, for each bidder i and each location l, a *synergy* term $w(i, l, l')$ which is the added value to bidder i in location l from having both locations l and l'. Therefore, if a bidder wins both locations l and l', his net valuation for the bundle (l, l') is $v(i, l) + v(i, l') + w(i, l, l') + w(i, l', l)$. In general, if a bidder wins a set of locations L', his valuation is $\sum_{l \in L'} v(i, l) + \sum_{l \in L'} \sum_{l' \in L'} w(i, l, l')$. We therefore associate a variable $y(i, l, l')$ indicating that bidder i has won both locations l and l', and can realize the associated synergy. For the formulation to work, however, we require that $w(i, l, l') \geq 0$ for all i, l and l'. That is, there should be no disadvantage to winning two locations (locations are complements). The BCA formulation is now as follows:

$$\max \sum_{i\in\mathcal{N},l\in\mathcal{L}} v(i,l)x(i,l) \quad + \sum_{i\in\mathcal{N},l\in\mathcal{L},l'\in\mathcal{L}} w(i,l,l')y(i,l,l') \tag{12}$$

$$\sum_{i\in\mathcal{N}} x(i,l) \;=\; 1 \qquad \forall l \in \mathcal{L} \tag{13}$$

$$y(i,l,l') \;\leq\; x(i,l) \qquad \forall i \in \mathcal{N},\ l,l' \in \mathcal{L} \tag{14}$$

$$y(i,l,l') \;\leq\; x(i,l') \qquad \forall i \in \mathcal{N},\ l,l' \in \mathcal{L} \tag{15}$$

$$x(i,l), y(i,l,l') \;\in\; \{0,1\} \qquad \forall i \in \mathcal{N},\ \forall l,l' \in \mathcal{L} \tag{16}$$

Constraints (14) and (15) prevent the program from realizing the synergies unless both locations are in fact awarded to the same bidder. Since this is a maximization program and we require $w(i,l,l') \geq 0$, the program will always realize the synergies whenever two synergistic locations are allocated to the same bidder.

Remark 1 *The program (12-16) is an instance of the forward-auction version of* BCA.

Proof: The value function $\tilde{c}$ is simply the concatenation of the vectors v and w, while the constraints (14-15) show the construction of the matrices A_i and b_i. Each bidder's type vector θ_i is simply the concatenation of the vectors v and w, capturing the valuation of any bundle of bids. $\square$

Locational synergies in general may involve much more than pair-wise synergistic terms. For example, the *number of contiguous blocks* awarded may have an impact on the value of the total bundle, as may other requirements such as the proportion of blocks which are large metropolitan areas (as opposed to rural areas), whether or not an entire state is covered, etc. We defer these generalizations to future research.

5.1.2 Bandwidth and Other Synergies

Bandwidth synergies occur when the same bidder in the same location is awarded two adjacent frequency bands, enabling that bidder to utilize the dead zone band between the two allocated bands at no further cost. A simple extension of the above model allows us to incorporate these bandwidth synergies within the BCA framework.

While the above provides a framework for two major synergies, one can conceive of several other synergistic effects which dictate the value of a bundle of several location-bandwidth pairs. A comprehensive study of incorporating all such synergies into the BCA framework is left open for future research.

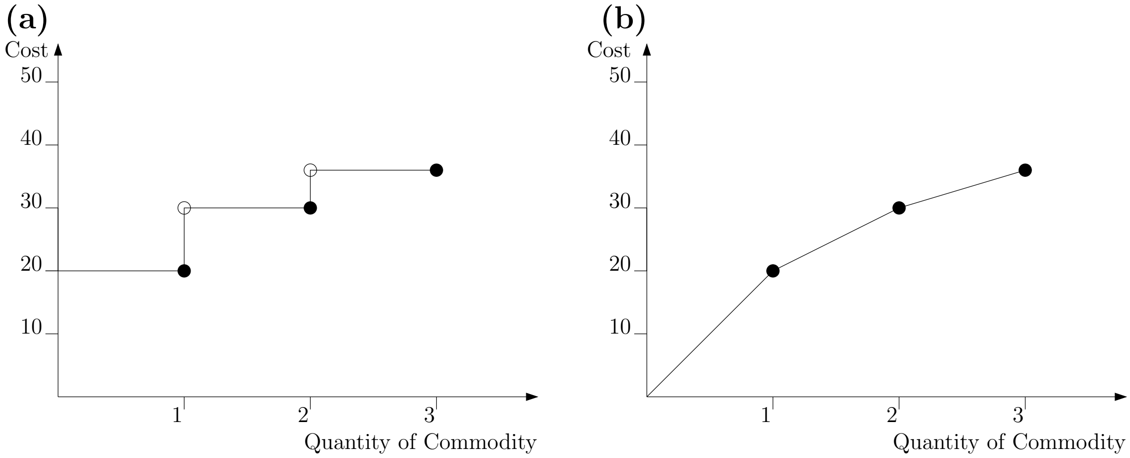

Figure 5: Examples of non-linear cost functions.

5.2 Commodities with Non-Linear Value Functions

5.2.1 Combinatorial Auctions for Single Commodities

Combinatorial auctions are typically thought of as auctions in which many different items are being auctioned simultaneously and the cost of providing a group of these items is not simply the sum of the individual items' costs. It is also possible, however, to think of combinatorial auctions in the context of a single commodity, in the case where the cost associated with that commodity is not linearly dependent on the quantity provided. For example, consider the cost function in Figure 5(a). If the bidder provides one unit the cost is twenty, but if two units are provided, the cost is only thirty (i.e. fifteen per unit), and when the provision increases to three units, the cost is thirty-five (i.e. $11\frac{2}{3}$ per unit). Thus, if bidders are only allowed to submit a single per-unit bid cost, they may over- or under- bid, depending upon the quantity of the commodity that they are awarded.

Conversely, the non-linearity of a cost function can be captured (or closely approximated) by a combinatorial auction, in which bidders are allowed to place multiple bids. Here, a "bundle" is really an upper and lower bound within which the commodity's cost is either fixed or linearly dependent on the quantity. The cost structure may be a step function (e.g., Figure 5(a)), in which case the bid for any quantity within a given range is constant. The bid is twenty if the quantity is in the range $[0, 1]$, thirty for the range $(1, 2]$, and thirty-five for the range $(2, 3]$. Alternatively, for the cost function in Figure 5(b), a bid with no fixed base cost and a marginal cost of twenty per unit would be placed for range $[0, 1]$, a bid with fixed base cost of ten plus marginal cost ten per unit would be placed for range $[1, 2]$, etc.

5.2.2 A Math-Programming Formulation for Single-Commodity Auctions

Single-commodity auctions with non-linear cost functions, such as those described in the preceding section, are a natural fit for BCA. The work on energy auctions of Hobbs et al. (2000) is an example of formulating a single-commodity auction with non-linear cost functions as a combinatorial auction. In this section, we limit our presentation to formulating the WDP. Our formulation is based on the use of binary variables to limit each bidder to a specific range. Once that range is specified, the cost function becomes linearized. We start by presenting the formulation for those cases where the value is constant within a given range, then extend this model to the case where the value is linear relative to the award within a given range.

Let $\mathcal{N}$ be the set of bidders, each of whom specifies a set of ranges R^i. Each range $r \in R^i$ is defined by an upper bound u^i_r, a lower bound l^i_r, and a constant cost for that range c^i_r. Let Q be the quantity of the commodity to be auctioned. For each bidder i and range $r \in R^i$, we define the binary variable z^i_r that takes value one if the quantity awarded to bidder i falls within the range r and zero otherwise. For each bidder i and range $r \in R^i$, we also define the continuous, non-negative variable y^i_r, which represents the quantity of the commodity *within range* r^i awarded to bidder i. If the award does not fall within range r^i, then y^i_r is zero. The formulation is then:

$$\min \sum_{i \in \mathcal{N}} \sum_{r \in R^i} c^i_r z^i_r \tag{17}$$

$$\text{subject to: } \sum_{r \in R^i} z^i_r = 1 \qquad \forall i \in \mathcal{N} \tag{18}$$

$$u^i_r z^i_r \geq y^i_r \geq l^i_r z^i_r \qquad \forall i \in \mathcal{N}, r \in R^i \tag{19}$$

$$\sum_{i \in \mathcal{N}} \sum_{r \in R^i} y^i_r = Q \tag{20}$$

$$z^i_r \in \{0,1\} \qquad \forall i \in \mathcal{N}, r \in R^i \tag{21}$$

$$y^i_r \geq 0 \qquad \forall i \in \mathcal{N}, r \in R^i. \tag{22}$$

The first set of constraints ensures that each bidder's award is within a single range; the objective function can be computed from these ranges, as each range corresponds to a constant fixed cost. The second and third set of constraints force the bidder's award to be within his/her designated range. The fourth constraint allocates the total quantity of the commodity across the set of bidders. The final constraints ensure the integrality, where appropriate, and non-negativity of the variables.

To accommodate linear cost functions within a given range, the bidders instead specify for each range $r \in R^i$ a fixed base cost f^i_r and a marginal per-unit cost m^i_r. The constraints above continue to define the feasible region of the problem. The objective function is modified to:

$$\min \sum_{i \in \mathcal{N}} \sum_{r \in R^i} (f_r^i z_r^i + m_r^i y_r^i). \tag{23}$$

This captures both the fixed cost associated with the designated range and the marginal cost associated with the size of the award.

Remark 2 *The programs (17-22) and (23, 18-22) are instances of* BCA.

Note that here we use Q in lieu of M since we are procuring multiple items of the same object, rather than multiple objects.

5.2.3 Energy Auctions

Energy auctions provide a number of examples of single commodity auctions where the cost has a non-linear dependence on the quantity. Such auctions can be single-seller/multiple buyer, single-buyer/multiple seller, or multiple-buyer/multiple-seller. Examples of all three cases and further references can be found in Hobbs et al. (2000). In particular, energy production is often characterized by alternating fixed costs and linear marginal costs. As certain thresholds are reached, it may become necessary to utilize an additional generator, with a fixed start-up cost. Such cost functions provide natural breakpoints for linearization as shown above.

5.3 Procurement Auctions with Capacity Constrained Suppliers

Purchasing agents often use procurement auctions to lower prices paid to suppliers, particulary when the item or contract is standard or well-specified and multiple suppliers can compete. When suppliers' production costs include significant fixed costs, a supplier's lowest price usually coincides with excess production capacity. Realizing that supplier bid prices are often directly tied to capacity constraints, it is important whether or not suppliers are able to estimate capacity. As we discuss below, this estimation problem is easily handled in the BCA framework.

Due to the capacity constraints, a supplier who offers to produce item j must anticipate how the offered production of j will affect his ability to produce items other than j. Capacity forecasts are required if items are auctioned off sequentially, since when bidding on items auctioned earliest the supplier must try to anticipate what his capacity position will be during subsequent auctions. This creates an exposure problem for the supplier, as committing his capacity through aggressive bidding in early auctions might backfire with lost opportunities for greater profit in later auctions (Elmaghraby 2003, Gallien and Wein 2005). This exposure problem can be avoided by auctioning off all contracts simultaneously.

Consider a setting in which a buyer seeks to purchase items in the set $\mathcal{M}$, where d_j is the buyer's demand for item $j \in \mathcal{M}$. We let $\mathcal{N}$ be the set of suppliers, and to keep the exposition concise we assume that for each supplier $i \in \mathcal{N}$, production quantity vector $x_i = (x_{i1}, x_{i2}, \ldots, x_{iM})$ costs the supplier $c_i^T x_i$ to produce and must satisfy the linear constraints $A_i x_i \leq b_i$, where $(a_{lj})_i \in A_i$ represents the amount of resource l consumed by one unit of item j and b_l represents the total amount of resource l available. For instance, the resources could be labor hours, machine time availability, or possibly production emission limits.

A canonical way to conduct this auction would be the CCA auction described in §4.1. While it removes the need for suppliers to make detailed, complex estimates of their capacity exposure problem (by auctioning all contracts simultaneously), its enumerative approach means every vector, and associated cost, for the feasible region must be communicated explicitly to the buyer by every supplier i.

Continuing with our theme of a bid-generation approach alternative, we propose that in lieu of enumerative bidding, the buyer simply ask each supplier i to reveal (c_i, A_i, b_i), a much simpler task. Then, the buyer's winner determination problem is just a linear program as follows, which is much easier to solve than the enumerative combinatorial problem.

$$\min \sum_{i \in \mathcal{N}} c_i x_i \tag{24}$$

$$\begin{aligned} \text{subject to: } A_i x_i &\leq b_i && \forall i \in \mathcal{N} && (25) \\ \sum_{i \in \mathcal{N}} x_{ij} &\geq d_j && \forall j \in \mathcal{M} && (26) \\ x_i &\geq 0 && \forall i \in \mathcal{N}. && (27) \end{aligned}$$

Remark 3 *The program (24-27) is an instance of* BCA.

While for brevity we have assumed linear capacity constraints and costs, a richer cost and capacity structure (e.g., that described in §5.2) could be incorporated with an analogous approach, which we leave for future research.

6 Conclusions and Future Research

In this paper, we address two critical limitations of traditional combinatorial auctions: the cognitive burden on bidders wrought by enumerative bidding, and the computational burden imposed on the auctioneer by an exponentially large winner determination problem. We address both problems simultaneously by a novel approach in which each bidder's bid-generating function is incorporated explicitly into the winner determination problem, rather than through enumerative bidding as in traditional models. This greatly reduces the cognitive burden on bidders, who simply communicate to the auctioneer parameters of their bid-generating functions. When the structure of the underlying bid-generating function is amenable, the resulting winner determination problem becomes tractable as well.

We focus on a key example where this is the case: procurement of truckload services. We show that the underlying bid-generating function is a minimum cost flow problem, and thus the winner determination problem can be posed as a simple multicommodity flow problem, which is known to be tractable in practice. It therefore becomes possible to solve instances of real-world size to optimality; that is, users can fully achieve economies of scope.

Future research in this application area includes additional operational considerations (some of which were discussed in §3.7), such as fleet capacities, round-trip limits, min/max restrictions on the number of carriers, etc. The question of how to incorporate into this model uncertainties in cost parameters (arising from the variability of the spot market for future loads, fuel costs, etc.) is also an interesting and important open research direction.

In addition to truckload procurement, we also introduce a number of other applications where the bid-generating functions appear to be amenable to our approach. These include wireless spectrum auctions, energy auctions, and procurement auctions with capacity-constrained suppliers. This opens a wide field of potential future research directions – adding more realism to the applications discussed in this paper and identifying new application areas – which will benefit from further collaboration between the mathematical programming and auction communities as well as application experts.

7 Acknowledgments

The authors would like to thank the Tauber Institute for Global Operations (formerly Tauber Manufacturing Institute) at the University of Michigan for initiating this project via their summer team internship program, and their continued support.

A Appendix: Proofs of Propositions and Theorems

Proposition 1 *For any problem of the form (5)-(10), replacing the integrality constraint (10) with the non-negativity constraint* $y^i_{jk} \geq 0 \;\; \forall i \in \mathcal{N}, j, k \in \mathcal{M}$ *results in a problem with the same optimal solution as the original problem.*

Proof: Consider any vector $\widehat{x}^i_j$ in which each element has integer value (i.e. either 0 or 1). Given such a vector, the problem then reduces to:

$$z^*(\mathcal{N}) = \min \sum_{i \in \mathcal{N}} \sum_{j \in \mathcal{M}} \sum_{k \in \mathcal{M}} e^i_{jk} y^i_{jk} \tag{28}$$

$$\text{subject to} \quad \sum_{k \in \mathcal{M}} y^i_{kj} = \widehat{x}^i_j \qquad \forall i \in \mathcal{N}, j \in \mathcal{M} \tag{29}$$

$$-\sum_{k \in \mathcal{M}} y^i_{jk} = -\widehat{x}^i_j \qquad \forall i \in \mathcal{N}, j \in \mathcal{M} \tag{30}$$

$$y^i_{jk} \in \{0, 1\} \qquad \forall i \in \mathcal{N}, j, k \in \mathcal{M}. \tag{31}$$

The left-hand side A matrix of this problem is totally unimodular, because each column contains a single element with value 1, a single element with value -1, and all other elements have value 0. Furthermore, the right-hand side vector is integer by supposition. Therefore, all extreme points of this polyhedron are integral. Thus, for any integer solution x, the values of y will be integer as well and we can therefore relax the integrality restriction. □

Proposition 4 *The cost function defined by (1)-(4) in the trucking procurement problem is sub-additive; that is, for any carrier i and sets of loads $s^i_{k_1}$ and $s^i_{k_2} \subseteq \mathcal{M} \setminus s^i_{k_1}$, we have $p^i_K \leq p^i_{k_1} + p^i_{k_2}$, where $s^i_K = s^i_{k_1} \cup s^i_{k_2}$.*

Proof: Let $y^*_{k_1}$ be the optimal solution to the program (1)-(4) instantiated for computing $p^i_{k_1}$, with $y^*_{k_2}$ defined similarly. Now consider the program (1)-(4) instantiated for computing p^i_K. The x terms in the program are defined by construction – an x variable is set to 1 if the corresponding load belongs to s^i_K. Define the vector y as follows: y_{jh} is set to 1 if and only if (i) either $j, h \in s^i_{k_1}$ and $y^*_{k1;j,h} = 1$, or (ii) $j, h \in s^i_{k_2}$ and $y^*_{k2;j,h} = 1$. In all other cases, set $y_{jh} = 0$. This defines a feasible solution, since $y^*_{k_1}$ and $y^*_{k_2}$ are feasible for the programs computing $p^i_{k_1}$ and $p^i_{k_2}$ respectively. The cost of this solution is $\hat{p}^i_K = p^i_{k_1} + p^i_{k_2}$. Since the optimal solution to the program has cost $p^i_K \leq \hat{p}^i_K$, the proposition holds. □

References

Ahuja, R., T. Magnanti, J. Orlin. 1993. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ.

An, N., W. Elmaghraby, P. Keskinocak. 2005. Bidding strategies and their impact on revenues in combinatorial auctions. *J. Revenue and Pricing Mgmt.* **3**(4) 337–357.

Anandalingam, G., R. Day, S. Raghavan. 2005. The landscape of electronic market design. *Management Sci.* **51**(3) 316–327.

Ausubel, L., P. Cramton, P. Milgrom. 2006. The clock-proxy auction: A practical combinatorial auction design. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Ausubel, L., P. Milgrom. 2006. The lovely and lonely Vickrey auction. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Ball, M., G. Donohue, K. Hoffman. 2006. Auctions for the safe, efficient, and equitable allocation of airspace system resources. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Bichler, M., J. Kalagnanam. 2005. Configurable offers and winner determination in multi-attribute auctions. *Eur. J. Oper. Res.* **160** 380–394.

Binmore, K., P. Klemperer. 2002. The biggest auction ever: The sale of the British 3G telecom licenses. *Econom. J.* **112** C74–C96.

Boutilier, C. 2002. Solving concisely expressed combinatorial auction problems. *Proc. 18th Natl. Conf. Artificial Intelligence*. 359–366.

Bureau of Economic Analysis. 2005. BEA News, U.S. Department of Commerce, news release August 31, 2005. http://www.bea.gov/bea/rels.htm.

Bureau of the Census. 2005. Statistics for Industry Groups and Industries: 2004, Annual Survey of Manufacturers. http://www.census.gov/prod/2005pubs/am0431gs1.pdf. M04(AS)-1.

Caplice, C. 1996. Optimization-based bidding, a new framework for shipper-carrier relationships. Ph.D. thesis, Massachusetts Institute of Technology.

Caplice, C., Y. Sheffi. 2006. Combinatorial auctions for truckload transportation. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Che, Y.-K. 1993. Design competition through multidimensional auctions. *RAND J. Econ.* **24**(4) 668–680.

Chen, R., R. Roundy, R. Zhang, G. Jamakiraman. 2005. Efficient auction mechanisms for supply chain procurement. *Management Sci.* **51**(3) 467–284.

Chopra, S., P. Meindl. 2007. *Supply Chain Management: Strategy, Planning and Operation*. Pearson Prentice Hall.

Clarke, E. H. 1971. Multipart pricing of public goods. *Public Choice* **11**(1) 17–33.

Cramton, P. 1997. The FCC spectrum auctions: An early assessment. *J. Econ. Mgt. Strategy* **6**(3) 431–495.

Cramton, P. 2002. Spectrum auctions. M. Cave, S. Majumdar, I. Vogelsgang, eds., *Handbook of Telecommunication Economics*. Elsevier.

Cramton, P., R. Steinberg, Y. Shoham. 2006. *Combinatorial Auctions*. MIT Press.

de Vries, S., R. Vohra. 2003. Combinatorial auctions: A survey. *INFORMS J. Comput.* **15**(3) 284–309.

Elmaghraby, W.J. 2003. The importance of ordering in sequential auctions. *Management Sci.* **49**(5) 673–682.

Epstein, R., L. Henríquez, J. Catalán, G. Weintraub, C. Martínez. 2002. A combinatorial auction improves school meals in Chile. *Interfaces* **32**(6) 1–14.

Federal Communications Commission. 2006. Auction of advanced wireless services licenses closes. Report No. AUC-06-66-F.

Figliozzi, M., H. Mahmassani, P. Jaillet. 2003. A framework for the study of carrier strategies in an auction based transportation marketplace. *TRB Annual Meeting*.

Franklin, M., M. Reiter. 1996. The design and implementation of a secure auction service. *IEEE Trans. Soft. Eng.* **22**(5) 302–312.

Gallien, J., L.M. Wein. 2005. A smart market for industrial procurement with capacity constraints. *Management Sci.* **51**(1) 76–91.

General Electric Corporation. 2001. Letter to share owners. *GE Annual Report 2000*.

Groves, T. 1973. Incentives in teams. *Econometrica* **41**(4) 617–631.

Günlük, O., L. Ladányi, S. de Vries. 2005. A branch-and-price algorithm and new test problems for spectrum auctions. *Management Sci.* **51**(3) 391–406.

Hobbs, B., M. Rothkopf, L. Hyde, R. O'Neill. 2000. Evaluation of a truthful revelation auction in the context of energy markets with nonconcave benefits. *J. Regulatory Econ.* **18**(1) 5–32.

Hohner, G., J. Rich, E. Ng, G. Reid, A.J. Davenport, J.R. Kalagnanam, H.S. Lee, C. An. 2003. Combinatorial and quantity-discount procurement auctions benefit Mars Inc. and its suppliers. *Interfaces* **33**(1) 23–35.

Huff, A. 2006. Freight by wire. http://www.etrucker.com/apps/news/article.asp?id=12683. Retrieved May 2006.

Internet Truckstop. 2006. The Internet Truckstop. http://www.truckstop.com. Retrieved May 2006.

Klemperer, P. 2002. What really matters in auction design. *J. Econom. Perspectives* **16** 169–190.

Krishna, V. 2002. *Auction Theory*. Academic Press, Elsevier Science.

Krishna, V., M. Perry. 1998. Efficient mechanism design. Tech. rep., The Pennsylvania State University. Working Paper.

Kudo, M. 1998. Secure electronic sealed-bid auction protocol with public key cryptography. *IEICE Trans. Fund. Electron. Commun. Comput. Sci.* **E81-A**(1) 20–27.

Kwon, R., G. Anandalingam, L. Ungar. 2005. Iterative combinatorial auctions with bidder-determined combinations. *Management Sci.* **51**(3) 407–418.

Ledyard, J., M. Olson, D. Porter, J Swanson, D. Torma. 2002. The first use of a combined-value auction for transportation services. *Interfaces* **32**(5) 4–12.

Lee, H., V. Padmanabhan, S. Whang. 1997. The bullwhip effect in supply chains. *MIT Sloan Management Review* **38**(3) 93–102.

McCabe, K., S. Rassenti, V. Smith. 1991. Smart computer-assisted markets. *Science* **254**(5031) 534–538.

McMillan, J. 1994. Selling spectrum rights. *J. Econ. Perspectives* **8**(3) 145–162.

Müller, R. 2006. Tractable cases of the winner determination problem. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Nisan, N. 2000. Bidding and allocation in combinatorial auctions. *Proc. ACM. Conf. Electronic Commerce (EC-00)*. 1–12.

Nisan, N. 2006. Bidding languages for combinatorial auctions. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Nisan, N., I. Segal. 2006. The communication requirements of efficient allocations and supporting prices. *J. Econom. Theory* **129**(1) 192–224.

Parkes, D. 2006. Iterative combinatorial auctions. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Pekec, A., M. Rothkopf. 2003. Combinatorial auction design. *Management Sci.* **49**(11) 1485–1503.

Plummer, C. 2003. Bidder response to combinatorial auctions in truckload procurement. Master's thesis, Massachusetts Institute of Technology.

Rassenti, S., V. Smith, R. Bulfin. 1982. A combinatorial auction mechanism for airport time slot allocation. *Bell J. Econ.* **13**(2) 402–417.

Rothkopf, M., A. Pekeč, R. Harstad. 1998. Computationally manageable combinational auctions. *Management Sci.* **44**(8) 1131–1147.

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

Sandholm, T., C. Boutilier. 2006. Preference elicitation in combinatorial auctions. P. Cramton, R. Steinberg, Y. Shoham, eds., *Combinatorial Auctions*. MIT Press.

Sandholm, T., D. Levine, M. Concordia, P. Martyn, R. Hughes, J. Jacobs, D. Begg. 2006. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. *Interfaces* **36**(1) 55–68.

Sandholm, T., S. Suri, D. Gilpin, D. Levine. 2005. Cabob: A fast optimal algorithm for winner determination in combinatorial auctions. *Management Sci.* **51**(3) 374–390.

Sheffi, Y. 2004. Combinatorial auctions in the procurement of transportation services. *Interfaces* **34**(4) 245–252.

Song, J., A. Regan. 2003. Combinatorial auctions for transportation service procurement: The carrier perspective. *TRB Annual Meeting*.

Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. Finance* **16** 8–37.

Wu, D., P. Kleindorfer. 2005. Competitive options, supply contracting, and electronic markets. *Management Sci.* **51**(3) 452–466.