

An Overview of Combinatorial Auctions¹

PETER CRAMTON

University of Maryland

and

YOAV SHOHAM

Stanford University

and

RICHARD STEINBERG

University of Cambridge

1. INTRODUCTION

An auction is *combinatorial* when bidders can place bids on combinations of items, called “packages,” rather than just individual items. Computer scientists are interested in combinatorial auctions because they are concerned with the expressiveness of bidding languages, as well as the algorithmic aspects of the underlying combinatorial problem. The combinatorial problem has attracted attention from operations researchers, especially those working in combinatorial optimization and mathematical programming, who are fascinated by the idea of applying these tools to auctions. Auctions have been studied extensively by economists, of course.² Thus, the newly emerging field of combinatorial auctions lies at the intersection of computer science, operations research, and economics.

In this article, we present a brief introduction to combinatorial auctions, based on our book, *Combinatorial Auctions* (MIT Press, 2006), in which we look at combinatorial auctions from all three perspectives. Indeed, our contribution is to take an integrated and comprehensive approach. We have done this in three ways. First, we have defined what we see as the five major subfields comprising combinatorial auctions: mechanisms, bidding and efficiency, complexity and algorithmic considerations, testing and implementation, and applications. The book is accordingly divided into five sections, with chapters written by the foremost experts on each topic within each subfield. Second, we have defined for the first time a common lan-

¹This article is based on our introductory chapter to *Combinatorial Auctions*, Peter Cramton, Yoav Shoham, and Richard Steinberg (eds.), MIT Press, 2006.

²Operations researchers were also active contributors to the early work on auctions, see for example, Friedman (1955) and Rothkopf (1969). Indeed, most of the early work on auctions first appeared in operations research journals.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 2007 ACM 1529-3785/2007/0700-0001 \$5.00

guage for combinatorial auctions, which has been accepted by researchers from all three communities, and have made a serious effort to use these terms consistently throughout the book. The most common terms are defined in the book's glossary. Third, the chapters are all cross-referenced; thus, for example, a mention of the Vickrey auction in any later chapter will refer back to the first chapter in the book by Ausubel and Milgrom on "The Lovely but Lonely Vickrey Auction." The foreword to the book is penned by Vernon L. Smith, who founded the field, together Stephen J. Rassenti and Robert L. Bulfin in 1982, with the publication of their paper, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation."

The book took four years from conception to completion. We were fortunate to attract as contributors most of the major researchers working in combinatorial auctions. The book was published by MIT Press in 2006.

2. COMBINATORIAL AUCTIONS IN PRACTICE

There are numerous examples of combinatorial auctions in practice. As is typical of many fields, practice precedes theory. Simple combinatorial auctions have been used for many decades in, for example, estate auctions. A common procedure is to auction the individual items, and then, at the end, to accept bids for packages of items. If a package bid exceeds the sum of the individual bids for the items in the package, then the items are sold as a package. In our book we consider a variety of much more general combinatorial auctions, but the key ingredient is the same as in this simple case: each bidder can submit bids on packages of items.

Recently, combinatorial auctions have been employed in a variety of industries. For example, they have been used for truckload transportation, bus routes, and industrial procurement, and have been proposed for airport arrival and departure slots, as well as for allocating radio spectrum for wireless communications services. Combinatorial auctions for radio spectrum have been conducted in both the United States and Nigeria. Most recently, the United Kingdom has proposed combinatorial auctions for spectrum, based in large part on the clock-proxy auction discussed later. In each application, the compelling motivation for the use of a combinatorial auction is the presence of complementarities among the items which differ across bidders. For example, a trucker's cost of handling shipments in one lane depends on its loads in other lanes. Similarly, a mobile phone operator may value licenses in two adjacent cities more than the sum of the individual license values, since the operator's customers value roaming between cities.

3. BASIC AUCTION THEORY

Auction theory is among the most influential and widely studied topics in economics of the last forty years. Auctions ask and answer the most fundamental questions in economics: who should get the goods and at what prices? In answering these questions, auctions provide the micro-foundation of markets. Indeed, many modern markets are organized as auctions.

To understand the role of combinatorial auctions, it is useful to step back and think about auctions in general. Some auction types are familiar, such as the ascending-bid English auction used in many online consumer auctions, or the first-price sealed-bid auction used in many public procurements. More fundamentally,

auctions are distinguished not only by the rules of the auction, such as ascending versus sealed-bid, but by the auction environment. These combinatorial auctions can be studied in a wide range of auction environments. Important features, including the numbers of sellers and buyers, the number of items being traded, the preferences of the parties, and the form of the private information participants have about preferences, all determine the auction environment.

The benchmark environment is the private value model, introduced by Vickrey (1961), which is discussed here in detail by Ausubel and Milgrom in Chapter 1. In the private value model, each bidder has a value for each package of items and these values do not depend on the private information of the other bidders. Each bidder knows his values, but not the values of the other bidders. Vickrey's seminal paper, mentioned in his 1996 Nobel Prize in economics, introduced the independent private value model, demonstrated equilibrium bidding behavior in a first-price auction, and then showed that truthful bidding could be induced as a dominant strategy by modifying the pricing rule: let each bidder pay the social opportunity cost of his winnings, rather than his bid. Finally, he showed in an example what would later be proven generally as the revenue equivalence theorem: different auction mechanisms that result in the same allocation of goods yield the same revenue to the seller.

Thus, when auctioning a single item to n bidders, whose payoffs are linear in the bidder's valuation of the item and money ($u_i = v_i - p$, where u_i is bidder i 's utility, v_i is i 's the value of the item, and p is the price paid for the item) and where each value is drawn independently from the same probability distribution, both the first-price and second-price auction award the item to the bidder with the highest value and yield the seller the same expected revenue.

Most of the chapters in our book use Vickrey's private value model and many make use of the Vickrey pricing rule, at least as a benchmark for comparison with alternative mechanisms.

Wilson (1969) took auction theory in a new direction. He introduced the common value auction model, in which items have the same value to all bidders, but this value is uncertain and depends on the private information of all bidders. He derived the first analysis of equilibrium bidding with common values, demonstrating the importance of conditioning one's bid on the negative information winning implies, and thus avoiding what would later be called the winner's curse—the tendency for bidders, who do not understand that winning is bad news about one's estimate of value, to pay more than the item is worth.

Milgrom extended Wilson's early papers in several ways. Most importantly, he introduced an auction model with both private value and common value elements. The private value model of Vickrey and common value model of Wilson represent two extreme cases. These extreme models are useful in deriving strong theoretical results, but most practical auction environments have both private and common value elements. Milgrom (1981) showed the importance of the monotone likelihood ratio property in obtaining results in a realistic hybrid model.³ In particular the

³A probability density function f satisfies the monotone likelihood ratio property if the ratio $f(v|t)/f(v|s)$ is weakly increasing in v for all $t > s$. Typically, $f(v|s)$ is the probability density of a bidder's value v conditional on the signal s (an estimate of value). Intuitively, the

monotone likelihood ratio property, together with Wilson's assumption of conditional independence, means that (1) bidders use monotonic bidding strategies and (2) that a monotonic strategy satisfying the first-order condition constitutes an equilibrium.

Milgrom's model led to the affiliated values model (Milgrom and Weber 1982) in which a bidder's value depends directly on the private information of all the bidders. The critical condition here, closely related to the monotone likelihood ratio property in Milgrom (1981), is that the bidders' signals, typically estimates of value, are affiliated random variables. This amounts to the plausible condition that if one bidder has a high signal of value, it is more likely that the signals of the other bidders are high. The paper shows that Vickrey's revenue equivalence result no longer holds when we introduce a common value element. In particular, the revenues from the standard auction formats differ and can be ranked. Formats like ascending auctions, in which the price is linked to more affiliated private information yield higher revenues.

The early work of Vickrey, Wilson, and Milgrom was largely focused on an equilibrium analysis and comparison of standard auction formats. Myerson led the development of mechanism design theory, which enables the researcher to characterize equilibrium outcomes of all auction mechanisms, and identify optimal mechanisms—those mechanisms that maximize some objective, such as seller revenues. His first application was to auctions. Myerson (1981) determined the revenue-maximizing auction with risk-neutral bidders and independent private information. He also proved a general revenue equivalence theorem that says that revenues depend fundamentally on how the items are assigned—any two auction formats that lead to the same assignment of the items yield the same revenues to the seller.

The trick in Myerson's analysis was recognizing that any auction can be represented as a direct mechanism in which bidders simultaneously report their private information and then the mechanism determines assignments and payments based on the vector of reports. For any equilibrium of any auction game, there is an equivalent direct mechanism in which bidders truthfully report types and agree to participate. Hence, without loss of generality we can look at incentive compatible and individually rational mechanisms to understand properties of all auction games. Incentive compatibility respects the fact that the bidders have private information about their values; individual rationality respects the bidders' voluntary participation decision. This key idea is known as the revelation principle (Myerson 1979).

Myerson and Satterthwaite (1983) use this technique to prove the general impossibility of efficient bargaining when it is not common knowledge that gains from trade exist; that is, when it is not certain that a mutually beneficial agreement is possible. This same impossibility extends to auctions in which both sellers and buyers possess private information, although efficiency becomes possible when the traders jointly own the items (Cramton, Gibbons, and Klemperer 1987). Likewise, if the roles of buyer and seller are not fixed *ex ante*, but the traders may take on either role depending on price, then efficient mechanisms exist (Wilson 1993).

likelihood of high values increases with the estimate of value.

These early papers led to the rapid development of auction theory in the 1980s and 1990s. In addition, large empirical and experimental literatures have sprung from the theory. This work is summarized in a number of articles and books, for example, McAfee and McMillan (1987), Kagel and Roth (1995), Klemperer (2000, 2004), Krishna (2002), and Milgrom (2004).

4. WHY COMBINATORIAL AUCTIONS?

A shortcoming of most of the work mentioned above (Milgrom 2004 is an exception) is the failure to recognize that in many auction environments bidders care in complex ways about the packages of items they win. The advantage of combinatorial auctions (CAs) is that the bidder can more fully express his preferences. This is particularly important when items are *complements*. Items are complements when a set of items has greater utility than the sum of the utilities for the individual items (for example, a pair of shoes is worth more than the value of a left shoe alone plus the value of a right shoe alone). The auction designer also derives value from CAs. Allowing bidders more fully to express preferences often leads to improved economic efficiency (allocating the items to those who value them most) and greater auction revenues.

However, alongside their advantages, CAs raise a host of questions and challenges. Our book is devoted to discussing these questions, as well as the considerable progress made in answering them.

5. TYPES OF COMBINATORIAL AUCTIONS

Our book begins in Part I with a description and analysis of various combinatorial auction mechanisms.

The most famous combinatorial auction is the combinatorial generalization of the Vickrey auction already mentioned, the Vickrey-Clarke-Groves (VCG) mechanism. Ausubel and Milgrom (Chapter 1) explore the question of why the Vickrey auction with its appealing theoretical properties is seen so little in practice. In a VCG auction (also called a Vickrey auction), bidders report their valuations for all packages; items are allocated efficiently to maximize total value. Each winner pays the opportunity cost of his winnings: the incremental value that would be derived by assigning the bidder's items according to their next best use among the other bidders. In this way, a winning bidder achieves a profit equal to his incremental contribution to total value, and it is a dominant strategy for the bidder to truthfully report his values. Achieving efficiency in truth-dominant strategies is remarkable. Nonetheless, there are serious shortcomings. Most importantly, bidders are asked to express values for all packages without the aid of any information about prices. Also, when complementarities are large, seller revenues can be too low,⁴ and adding bidders or increasing bidder values can reduce seller revenue.

In Chapter 2, Parkes examines iterative combinatorial auctions. A major motivation for an iterative process is to help the bidders express their preferences by providing provisional pricing and allocation information. This information helps the bidders focus their valuation efforts on options that are most relevant.

⁴Goods are substitutes when increasing the price of one does not reduce demand for the other.

In Chapter 3, Ausubel and Milgrom consider the ascending proxy auction (Ausubel and Milgrom 2002) as an alternative to the Vickrey auction. Each bidder submits valuation information to a proxy agent. The proxy agents bid iteratively, bidding on the most profitable package, whenever the proxy agent is not a provisional winner. The auction ends when no proxy agent who is not a provisional winner has a profitable bid. The ascending proxy auction allows for bidders to have budget constraints. In the absence of budget constraints, and when goods are substitutes for all bidders, the ascending proxy auction yields the same outcome as the Vickrey auction. More generally, the ascending proxy auction finds a bidder-optimal point in the core with respect to the reported preferences. Moreover, all bidder-optimal core points are Nash equilibria in the auction game, if we assume full information about values (each bidder knows the values of the other bidders). The ascending proxy auction addresses many of the drawbacks of the Vickrey auction in environments with some complements.

The simultaneous ascending auction (SAA) is studied by Cramton in Chapter 4. The SAA is not a combinatorial auction, since bids in a SAA are placed for individual items, rather than packages of items. Yet the SAA has proven to be a highly effective method of auctioning many related items (see Cramton 1998, 2002 and Milgrom 2004). Simultaneous sale and ascending bids enables price discovery, which helps bidders build desirable packages of items. The SAA remains a useful benchmark for comparison with true combinatorial auctions.

In Chapter 5, Ausubel, Cramton, and Milgrom propose the clock-proxy auction as a practical combinatorial design. A clock auction is followed by a best-and-final proxy round. The clock auction is a simple dynamic auction. In each round, the auctioneer announces a price. Each bidder then indicates the quantity it desires to buy at that price. In subsequent rounds, the price increases, and each bidder again expresses the quantity it desires to buy at the new price. This process is repeated until there is no excess demand. The clock-proxy approach combines the simple and transparent price discovery of the clock auction with the efficiency of the proxy auction. Linear pricing is maintained as long as possible, but then is abandoned in the proxy round to improve efficiency and enhance seller revenues. The approach has many advantages over the simultaneous ascending auction. In particular, the clock-proxy auction has no exposure problem, eliminates incentives for demand reduction, and prevents most collusive bidding strategies. Without the best-and-final proxy round, the authors present an iterative combinatorial auction that can be implemented as a simple clock auction, avoiding all computational complexity issues in a process with highly useful price discovery (Ausubel and Cramton 2004). This auction format recently has been used in over two dozen high-stake auctions in several countries and several industries.

Chapter 6 discusses a combinatorial auction procedure called PAUSE, proposed by Frank Kelly and Richard Steinberg, which relieves the auctioneer of having to face the “winner determination problem” which as discussed below is a computationally intractable problem. Under PAUSE, the burden of evaluating a combinatorial bid is transferred to the bidder making the bid; the auctioneer need only confirm the bid’s validity, a computationally tractable problem. As a consequence, although PAUSE permits all combinatorial bids, the procedure is both compu-

tationally tractable for the auctioneer and transparent to the bidders. In their chapter, Land, Powell, and Steinberg focus specifically on bidder behavior under PAUSE.

6. BIDDING AND EFFICIENCY

As mentioned above, combinatorial auctions give rise to a host of interesting questions and challenges. To begin with, there is the question of what should be the bidding language—how bidder preferences are expressed. Different choices vary in expressiveness and in simplicity. A bid in an auction is an expression of the bidder’s preference for various outcomes. The most direct way of capturing such a preference is to have a bidder attach a monetary value to each possible allocation. This allows one to express all possible preferences, but it is not simple. Given n bidders and m items, it requires a bidder to submit a bid of size n^m . If we assume no externalities, so that each bidder cares only about the items he himself receives, the complexity drops to 2^m , which is still impractical for all but small m .

Part II of our book addresses both bidding languages and questions of efficiency. Auction theory generally assumes a fixed number of bidders with each bidder acting independently according to the rules of the auction. One simple deviation from this model is for a single bidder to act as multiple bidders. Such pseudonymous bidding is the subject of Chapter 7. Yokoo shows that the Vickrey auction is not immune to this problem, unless a bidder submodularity condition is satisfied. And indeed all efficient auctions suffer from this problem. It sometimes is profitable for a bidder to bid as multiple bidders, rather than one, and this undermines efficiency.

In Chapter 8, Bikhchandani and Ostroy examine the connection between efficient auctions for many items, and duality theory. The Vickrey auction can be thought of as an efficient pricing equilibrium, which corresponds to the optimal solution of a particular linear programming (LP) problem and its dual. A “buyers are substitutes” condition is necessary and sufficient for the pricing equilibrium to yield the Vickrey outcome. Thus, when buyers are substitutes, an efficient pricing equilibrium can be obtained with any LP algorithm. The simplex algorithm can be thought of as a static approach to determining the Vickrey outcome. Alternatively, the primal-dual algorithm can be thought of as a decentralized and dynamic method of determine the pricing equilibrium, as in the ascending proxy auction of Chapter 3.

In Chapter 9, Nisan examines a variety of bidding languages and their properties. For example, we see there that OR (“additive-or”) bids, which allow the bidder to make *non-exclusive* offers on bundles, can capture all, and only, the super-additive valuations. In contrast, XOR (“exclusive-or”) bids, which allow the bidder to make *exclusive* offers on bundles, can capture all valuations, though they may require an exponentially longer expression than the OR bids. However, asking an agent to disclose a full valuation function is often not necessary, since many parts of it might be irrelevant for computing the allocation.

In Chapter 10, Sandholm and Boutilier look at ways in which the valuation function of agents can be elicited piecemeal, as needed by the auctioneer. One of the questions there is what form the queries may take. Sandholm and Boutilier consider several primary forms; these include queries about absolute bundle values, differ-

ences between two bundle values, a simple ordering on bundle values, and several others. Among the experimental results they show is the fact that in practice only a small fraction of the preferences need to be revealed. Among the theoretical results presented are some natural valuation classes where preferences can be elicited with a polynomial number of queries even in the worst case; the fact that even if the real preferences only fall approximately into these classes, an approximation can be found with a polynomial number of queries; and the fact that there can be super-exponential power in interleaving queries across agents (i.e., deciding what to ask an agent based on what others have revealed).

Segal in Chapter 11 asks how many bits of information are required to compute an efficient allocation, regardless of the protocol used and disregarding issues of incentives. One result states that any mechanism that is guaranteed to compute an efficient allocation must necessarily also discover supporting prices (though these will in general be neither anonymous nor linear). The main question addressed by Segal is how one can trade off the extent of communication required with the economic surplus gained. For example, the trivial protocol in which bidders communicate their value for the entire set of goods, which is allocated to the highest bidder (again, ignoring the issue of incentives), guarantees $1/n$ of the available surplus (where n is the number of bidders) while requiring a single bid from each bidder. A more elaborate mechanism yields $1/\sqrt{m}$ of the available surplus, where m is the number of goods. Interestingly, this is also a lower bound for any protocol whose running time is polynomial in m .

7. COMPUTATIONAL COMPLEXITY AND ALGORITHMIC CONSIDERATIONS

Once the bidding language is fixed, the question remains as to how to compute the allocation, given a set of bids. This problem, called the *winner determination problem* (WDP) has received considerable attention in the literature, and is the primary focus of Part III.

In Chapter 12, Lehmann, Müller and Sandholm provide a precise formulation of the problem and explore its basic complexity properties. The problem is this: Given a set of bids in a combinatorial auction, find an allocation of items to bidders, including the possibility that the auctioneer retains some items, that maximizes the auctioneer's revenue. The problem, which is most naturally represented as an integer program (IP), is inherently complex. Specifically, it is NP-hard, meaning that there does not exist a polynomial-time algorithm that is guaranteed to compute the optimal allocation. Even worse, the problem is not uniformly approximable, in the following sense: there does not exist a polynomial-time algorithm and a constant d that, for all inputs, produces an answer that is at least $1/d$ of the correct optimal answer.

We then follow this sobering introduction to the WDP with some good news. First, in Chapter 13, Müller explores some constraints on the set of bids that ensure that a polynomial-time solution does exist. One such condition is for the constraint matrix to be *totally unimodular*. A special case of this is of *linear goods*; for example, if each bid is for some contiguous stretch of time on a shared machine, the problem can be solved in quadratic time. Then, in Chapter 14, Sandholm looks at algorithms for solving the general problem. While we know that in the worst

case any algorithm will run in exponential time, there exist rules of thumb for searching the space of allocations that in practice allow us to solve large problems (for example, with hundreds of thousands of bids and tens of thousands of items). Sandholm concentrates on *complete* heuristics, ones that guarantee that an optimal solution is found but do not guarantee the running time.

The discussion of the WDP in Chapters 12, 13, and 14 ignores issues of incentives. The optimization is assumed to be inherited from some mechanism, such as the VCG mechanism, but solved without regard to the originating mechanism. As discussed, these problems are computationally hard, and sometimes admit only sub-optimal solutions. In Chapter 15, Ronen looks at the impact of such sub-optimal optimization on the incentive properties of mechanisms. For example, it is shown that with sub-optimal procedures, the VCG mechanism is no longer individually rational, nor is it incentive compatible. However, a modification of VCG is presented that restores individual rationality and, to a certain extent, incentive compatibility. The chapter covers several other topics, including a non-VCG mechanism that is computationally easy and incentive compatible.

In the final chapter of Part III, Chapter 16, Peke and Rothkopf consider appropriate ways to reduce or avoid computational difficulties in combinatorial auctions. The authors briefly review the computational issues in combinatorial auction design, the context of auction design including the information available to the designer, and properties that the auction designer must trade off in selecting the auction format and procedures. The major part of their chapter discuss opportunities for mitigating computational problems at four points in the auction: before bid submission, at the time of bid submission, after bid submission but before the announcement of a tentative set of winning bids, and after the announcement of a tentative set of winning bids.

8. TESTING AND IMPLEMENTATION

In Part IV, we pick up the question of how to test experimentally the various proposed solutions to the WDP as well as how best to test and implement mechanisms from Part I.

In Chapter 17, Hoffman, Menon, van den Heever, and Wilson consider how best to implement the ascending proxy auction of Chapter 3. Three approaches for accelerating the algorithm are considered. The first involves working backward from the efficient allocation and starting with the Vickrey prices, which provide a lower bound on prices. The second approach, increment scaling, solves the problem with large bid increments and then backs up and solves the problem again with reduced increments until the desired accuracy is obtained. The third approach combines the previous two. These three approaches each dramatically reduce the number of iterations needed to determine the ascending proxy outcome.

In Chapter 18, Leyton-Brown and Shoham present the *Combinatorial Auction Test Suite* (CATS). CATS is a publicly available software package that generates a variety of winner determination problems. Specifically, it implements several parameterized families of bid distributions, some based on real-world applications (such as transportation networks), and some on historical distributions used by researchers in the field. The goal of CATS is to serve as a uniform test suite for

WDP algorithms, and it has been used widely in this capacity.

In Chapter 19, Leyton-Brown, Nudelman and Shoham use CATS to evaluate the running times of algorithms for the winner determination problem. The difficulty is that, since the problem is NP-complete, even the best heuristic procedures will take exponential time for some instances. In many applications, it is important to know in advance how long a given algorithm will run on a given instance (for example, an auction for energy production tomorrow needs to determine a schedule of operation well in advance of tomorrow). The authors describe how machine learning techniques can be used to predict this running time reliably, and which features of a given instance are most predictive of this running time. As a bonus, they also describe a portfolio approach to the WDP, whereby several competing algorithms are pressed into service, and for each instance the algorithm that is predicted to perform best is chosen.

9. FOUR IMPORTANT APPLICATIONS OF COMBINATORIAL AUCTIONS

In their seminal paper on combinatorial auctions, Rassenti, Smith, and Bulfin (1982) present a sealed-bid combinatorial auction for the allocation of airport time slots (i.e., takeoff and landing slots) to competing airlines.⁵ Even if landing slots are bundled with takeoff slots (much like left and right shoes are bundled), the need for a combinatorial auction follows from the differing ways that airlines value packages of slots: some are substitutes, some are complements, and the valuations vary across airlines. Congestion at many major airports is becoming an increasingly important problem. The Federal Aviation Administration is now evaluating a combinatorial auction approach for New York's LaGuardia airport.

The final section of our book, Part V, considers four important applications of combinatorial auctions. Chapter 20 takes up the topic of auctions for airport time slots. Ball, Donohue, and Hoffman provide suggestions for mechanisms for air transportation systems to both expand capacity and to assure that the current, limited capacity is used both safely and efficiently. The authors begin by providing a description of the history of the U.S. Aviation System, detail current procedures for allocating landing time slots, and explain how market-clearing mechanisms might be able to rectify many of the shortcomings of the current system. They include a presentation of some of the important elements that should be included in combinatorial auctions in this setting.

In Chapter 21, Caplice and Sheffi explore how combinatorial auctions are being used for the procurement of freight transportation services, focusing on those attributes of transportation that make combinatorial auctions especially attractive, as well as describing some of the unique elements of transportation auctions. They present such auctions first from the perspective of the auctioneer, i.e., the shipper, then from the perspective of the bidder, i.e., the carrier. This is followed by a

⁵This was the first major paper on combinatorial auctions. It introduced many important ideas, such as the mathematical programming formulation of the auctioneer's problem, the connection between the winner determination problem and the set packing problem as well as the concomitant issue of computational complexity, the use of techniques from experimental economics for testing combinatorial auctions, and consideration of issues of incentive compatibility and demand revelation in combinatorial auctions.

discussion of the relationships between shippers and carriers, since the contracts that govern them have certain characteristics that distinguish them to some extent from auctions for other applications discussed elsewhere in our book. In fact, the types of bids used in transportation are distinctive to that industry, so there is an entire section discussing them. In this industry, the winner determination problem is known as the “Carrier Assignment Problem,” which is discussed next in the chapter. Finally, the authors present lessons from practice.

In Chapter 22, we move from the private sector to the public sector. As Cantillon and Pesendorfer explain, the London bus routes market provides an early example of the use of a combinatorial auction format in public procurement. The authority responsible for the provision and procurement of public transport services in the Greater London area—valued at \$900 million—was London Regional Transport (LRT). The authors present the four major issues faced by LRT. First, what should be the set of contracts auctioned? Second, how should LRT auction these contracts? Third, who should be allowed to participate? Finally, which criteria should they use to award the contracts? The authors also discuss the motivations for submitting a package bid, a description of their data together with summary statistics, and finally their empirical analysis.

The final chapter of our book, Chapter 23, discusses combinatorial auctions for industrial procurement, which is potentially one of the largest application domains for combinatorial auctions. As pointed out by the authors, Bichler, Davenport, Hohner, and Kalagnanam, CAs have already turned into a topic of interest for software vendors and procurement managers in the business-to-business domain. However, despite reports of the existence of a number of applications of combinatorial auctions in industrial procurement, documentation and public information on design details are rare—possibly because of efforts to protect proprietary information. This chapter describes current practice in this domain, including a case study at Mars, Inc.

10. CONCLUSION

Both the research and practice of combinatorial auctions have grown rapidly over the past dozen years. Our book aims to make this knowledge accessible to a broad group of researchers and practitioners. However, we have a more ambitious goal. Our hope is that, by integrating the work from the three underlying disciplines of economics, operations research, and computer science, progress on combinatorial auctions will be enhanced. Our book lays the foundation by aggregating and harmonizing the research on combinatorial auctions. We are confident that the profession will continue to develop the theory and application of combinatorial auctions. We hope that the foundation provided in our book will help in this process.

11. REFERENCES

- Ausubel, Lawrence M. and Peter Cramton (2004), “Auctioning Many Divisible Goods,” *Journal of the European Economic Association*, 2, 480-493, April-May.
 Ausubel, Lawrence M. and Paul Milgrom (2002), “Ascending Auctions with Package Bidding,” *Frontiers of Theoretical Economics*, 1: 1-45,
www.bepress.com/bejte/frontiers/vol1/iss1/art1.

- Cramton, Peter (1998), "Ascending Auctions," *European Economic Review*, 42, 745-756.
- Cramton, Peter (2002), "Spectrum Auctions," in Martin Cave, Sumit Majumdar, and Ingo Vogelsang, eds., *Handbook of Telecommunications Economics*, Amsterdam: Elsevier Science B.V., Chapter 14, 605-639.
- Cramton, Peter, Robert Gibbons, and Paul Klemperer (1987), "Dissolving a Partnership Efficiently," *Econometrica*, 55, 615-632.
- Friedman, Lawrence (1955), "A Competitive Bidding Strategy," *Operations Research*, 4, 104-112.
- Kagel, John H. and Alvin E Roth (1995), *The Handbook of Experimental Economics*, Princeton University Press.
- Klemperer, Paul (2000), *The Economic Theory of Auctions*, Cheltenham, UK: Edward Elgar.
- Klemperer, Paul (2004) *Auctions: Theory and Practice*, Princeton, NJ: Princeton University Press.
- Krishna, Vijay (2002), *Auction Theory*, San Diego, CA: Academic Press.
- McAfee, R. Preston and John McMillan (1987), "Auctions and Bidding," *Journal of Economic Literature*, 25, 699-738.
- Milgrom, Paul (1981), "Rational Expectations, Information Acquisition, and Competitive Bidding," *Econometrica*, 49, 921-943.
- Milgrom, Paul (2004), *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.
- Milgrom, Paul and Robert J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50, 1089-1122.
- Myerson, Roger B. (1979), "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47, 61-73.
- Myerson, Roger B. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58-73.
- Rassenti, Stephen J., Vernon L. Smith, and Robert L. Bulfin (1982), "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell Journal of Economics*, 13, 402-417.
- Rothkopf, Michael H. (1969), "A Model of Rational Competitive Bidding," *Management Science*, 15, 362-372.
- Vickrey, William (1961), "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance*, 16, 8-37.
- Wilson, Robert (1969), "Competitive Bidding with Disparate Information," *Management Science*, 15, 446-448.
- Wilson, Robert (1993), "Design of Efficient Trading Procedures," in Daniel Friedman and John Rust (eds.), *The Double Auction Market: Institutions, Theories, and Evidence*, Reading, MA: Addison-Wesley Publishing Company, Chapter 5, 125-152.