AI magazine

Volume 28 • Number 3 Fall 2007

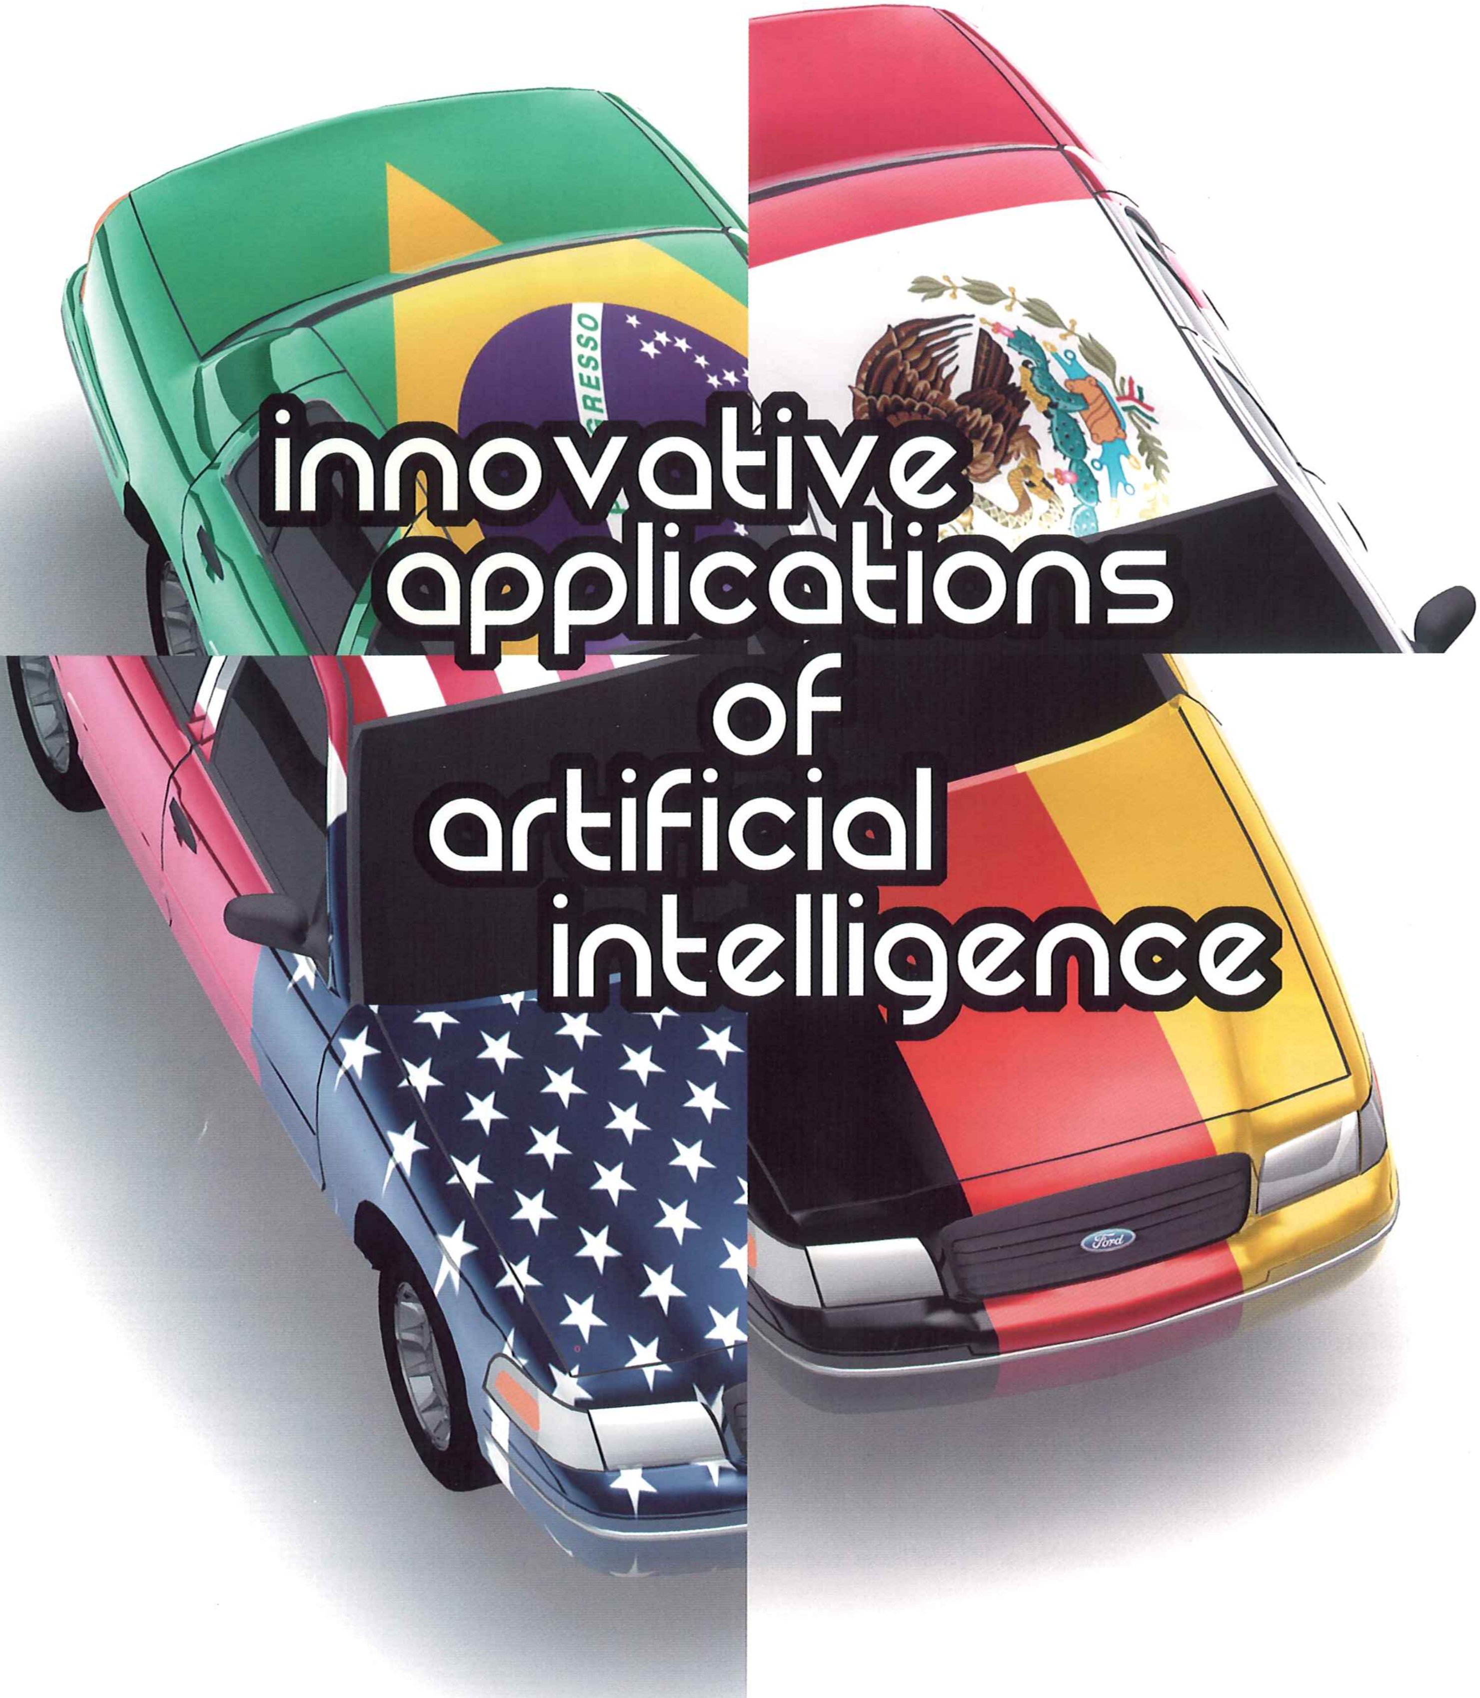

Expressive Commerce and Its Application to Sourcing: How We Conducted $35 Billion of Generalized Combinatorial Auctions

Tuomas Sandholm

■ Sourcing professionals buy several trillion dollars worth of goods and services yearly. We introduced a new paradigm called *expressive commerce* and applied it to sourcing. It combines the advantages of highly expressive human negotiation with the advantages of electronic reverse auctions. The idea is that supply and demand are expressed in drastically greater detail than in traditional electronic auctions and are algorithmically cleared. This creates a Pareto efficiency improvement in the allocation (a win-win between the buyer and the sellers), but the market-clearing problem is a highly complex combinatorial optimization problem. We developed the world's fastest tree search algorithms for solving it. We have hosted $35 billion of sourcing using the technology and created $4.4 billion of hard-dollar savings plus numerous harder-to-quantify benefits. The suppliers also benefited by being able to express production efficiencies and creativity, and through exposure problem removal. Supply networks were redesigned, with quantitative understanding of the trade-offs, and implemented in weeks instead of months.

Historical Backdrop on Sourcing

Sourcing, the process by which companies acquire goods and services for their operations, entails a complex interaction of prices, preferences, and constraints. The buyer's problem is to decide how to allocate the business across the suppliers.

Traditionally, sourcing decisions have been made through manual, in-person negotiations. The advantage is that there is a very expressive language for finding, and agreeing to, win-win solutions between the supplier and the buyer. The solutions are implementable because operational constraints can be expressed and taken into account. On the downside, the process is slow, unstructured, and nontransparent. Furthermore, sequentially negotiating with the suppliers is difficult and leads to suboptimal decisions. (This is because what the buyer should agree to with a supplier depends on what other suppliers would have been willing to agree to in later negotiations.) The one-to-one nature of the process also curtails competition.

These problems have been exacerbated by a dramatic shift from plant-based sourcing to global corporatewide (category-based rather than plant-based) sourcing since the mid-1990s. This transition is motivated by a corporation's desire to leverage its spending across plants in order to get better pricing and better understanding and control of the supply chain while at the same time improving supplier relationships. (See, for example, Smock [2004].) This transition has yielded significantly larger sourcing events that are inherently more complex.

During this transition, there has also been a shift to electronic sourcing where suppliers submit offers electronically to the buyer. The buyer then decides, using software, how to allocate the business. Advantages of this approach include speed of the process, structure and transparency, global competition, and simultaneous negotiation with all suppliers (which removes the difficulties associated with the speculation about later stages of the negotiation process, discussed above).

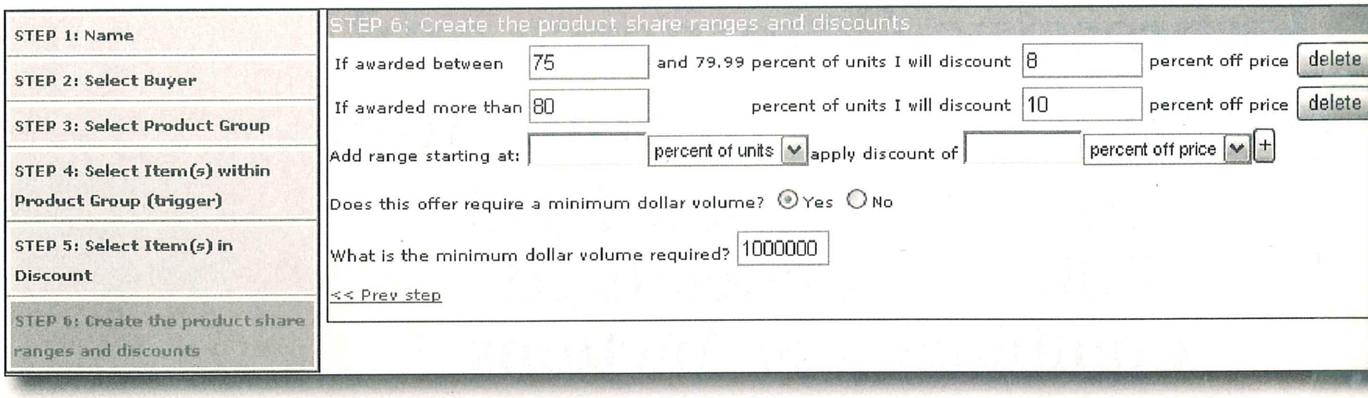

Figure 1. A Relatively Simple Discount Schedule.

This screenshot is from an actual sourcing event. The scope of the trigger of the discount (STEP 4) can be different than the scope of the items to which the discount is to be applied (STEP 5).

The most famous class of electronic sourcing systems—which became popular in the mid-1990s through vendors such as FreeMarkets (now part of Ariba), Frictionless Commerce (now part of SAP), and Procuri—is a *reverse auction*. The buyer groups the items into lots in advance and conducts an electronic descending-price auction for each lot. The lowest bidder wins. (In some cases "lowness" is not measured in terms of price but in terms of an ad hoc score that is a weighted function that takes into account the price and some nonprice attributes such as delivery time and reputation.)

Reverse auctions are not economically efficient; that is, they do not generally yield good allocation decisions. This is because the optimal bundling of the items depends on the suppliers' preferences (which arise, among other considerations, from the set, type, and time-varying state of their production resources), which the buyer does not know at the time of lotting. Lotting by the buyer also hinders the ability of small suppliers to compete. Furthermore, reverse auctions do not support side constraints, yielding two drastic deficiencies: (1) the buyer cannot express his or her business rules; thus the allocation of the auction is unimplementable and the "screen savings" of the auction do not materialize in reality, and (2) the suppliers cannot express their production efficiencies (or differentiation), and are exposed to bidding risks (for example, if the buyer's desired package consists of multiple lots). In short, reverse auctions assume away the complexity that is inherent in the problem and dumb down the events rather than embracing the complexity and viewing it as a driver of opportunity. It is therefore not surprising that there are strong broad-based signs that reverse auctions have fallen in disfavor.

The New Paradigm: *Expressive Commerce*

In 1997 it dawned on me that it is possible to achieve the advantages of both manual negotiation and electronic auctions while avoiding the disadvantages. The idea is to allow supply and demand to be expressed in drastically more detail (as in manual negotiation) while conducting the events in a structured electronic marketplace where the supply and demand are algorithmically matched (as in reverse auctions). The new paradigm, which we call *expressive commerce™* (or *expressive competition™*), was so promising that I decided to found CombineNet, Inc., to commercialize it.

The finer-grained matching of supply and demand yields Pareto improvements (that is, win-win solutions) between the buyer and the suppliers. However, matching the drastically more detailed supply and demand is an extremely complex combinatorial optimization problem. We developed the world's fastest algorithms for optimally solving it. These algorithms are incorporated into the market-clearing engine, *Clear Box™*, at the core of our flagship product, *the Advanced Sourcing Application Platform* (ASAP).

Expressive commerce has two sides: *expressive bidding™* and *expressive allocation evaluation* (also called *expressive bid taking™*) (Sandholm and Suri 2001).

Expressive Bidding

With expressive bidding, the suppliers can express their offers creatively, precisely, and conveniently using expressive and compact statements that are natural in the suppliers' business. Our expressive bidding takes on sev-

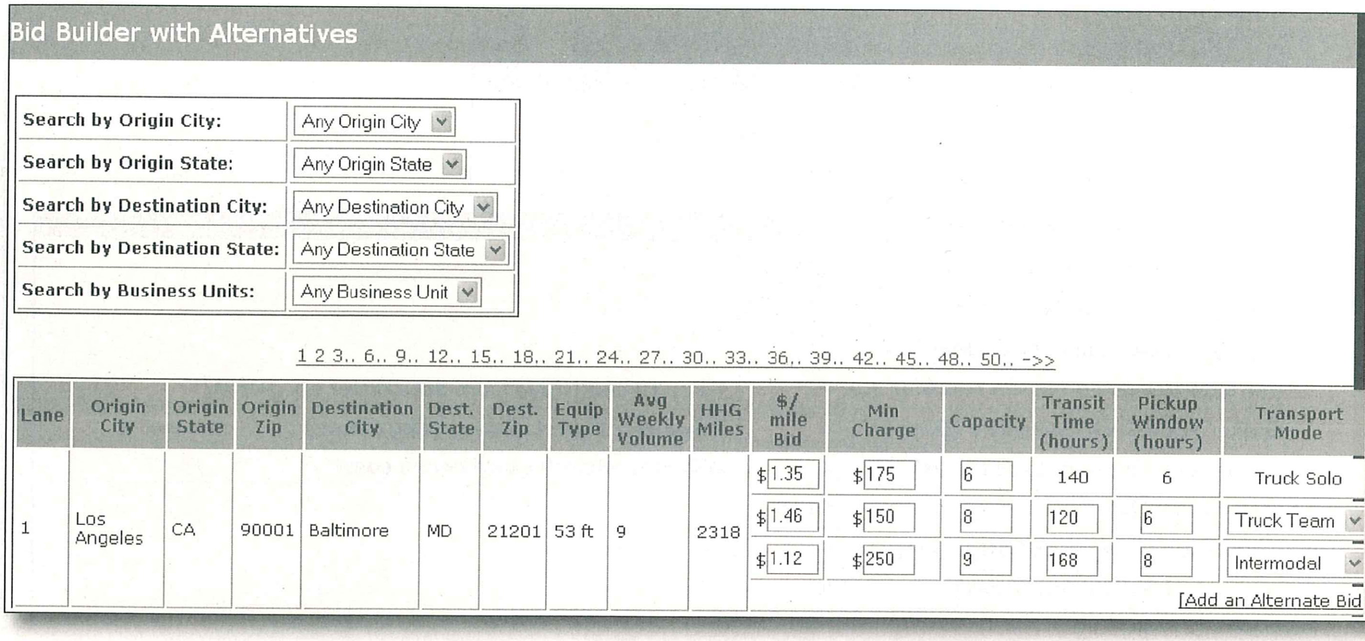

Figure 2. A Simple Example of Bidding with Alternates, Cost Drivers, Attributes, and Constraints.

This piece of screen is from an actual event for sourcing truckload transportation. The figure shows part of the bid by one bidder on one item.

eral forms. ASAP supports the following forms of expressive bidding, among others, all in the same event. These forms will be explained in the following paragraphs.

Bidding on an arbitrary number of self-constructed packages of items (rather than being restricted to bidding on predetermined lots as in basic reverse auctions). The packages can be expressed in more flexible and more usable forms than what is supported in vanilla combinatorial auctions. For example, the bidder can specify different prices on the items if the items get accepted in given proportions, and the bidder can specify ranges for these proportions, thus allowing a huge space of packages to be captured by one compact expression.

Conditional discount offers. The trigger conditions and the effects can be specified in highly flexible ways.

Rich forms of discount schedules. (Simpler forms of discount schedules have already been addressed in the literature (Sandholm and Suri 2001, Sandholm and Suri 2002, Hohner et al. 2003).) Figure 1 shows a simple example. Richer forms allow the bidder to submit multiple discount offers and to control whether and how they can be combined (for example, sequenced in particular ways). Discount triggers can be expressed as dollars or units, and as a percentage or an absolute.

A broad variety of side constraints—such as capacity constraints (Sandholm and Suri 2001).

Multiattribute bidding (Sandholm and Suri 2001). This allows the buyer to leave the item specification partially open, so the suppliers can pick values for the item attributes—such as material, color, and delivery date—in a way that matches their production efficiencies. This is one way in which the suppliers can also express alternate items.

Free-form expression of alternates. This fosters unconstrained creativity by the suppliers.

Expression of cost drivers. In many of our events, the buyer collects tens or hundreds of cost drivers (sometimes per item) from the suppliers. By expressing cost drivers, the bidder can concisely implicitly price huge numbers of items and alternates. Figures 2 and 3 illustrate bidding with attributes and cost drivers.

All of these expressive bidding features of ASAP have been extensively used by CombineNet's customers. ASAP supports bidding through web-based interfaces and through spreadsheets. In some cases, catalog prices from databases have also been used.

CombineNet's expressive bidding is flexible in the sense that different suppliers can bid in different ways, using different offer constructs. In fact, some suppliers may not be sophisticated enough to bid expressively at all, yet they can participate in the same sourcing events using traditional bidding constructs in the same system. This paves a smooth road for adoption, which does not assume sudden process changes at the participating organizations.

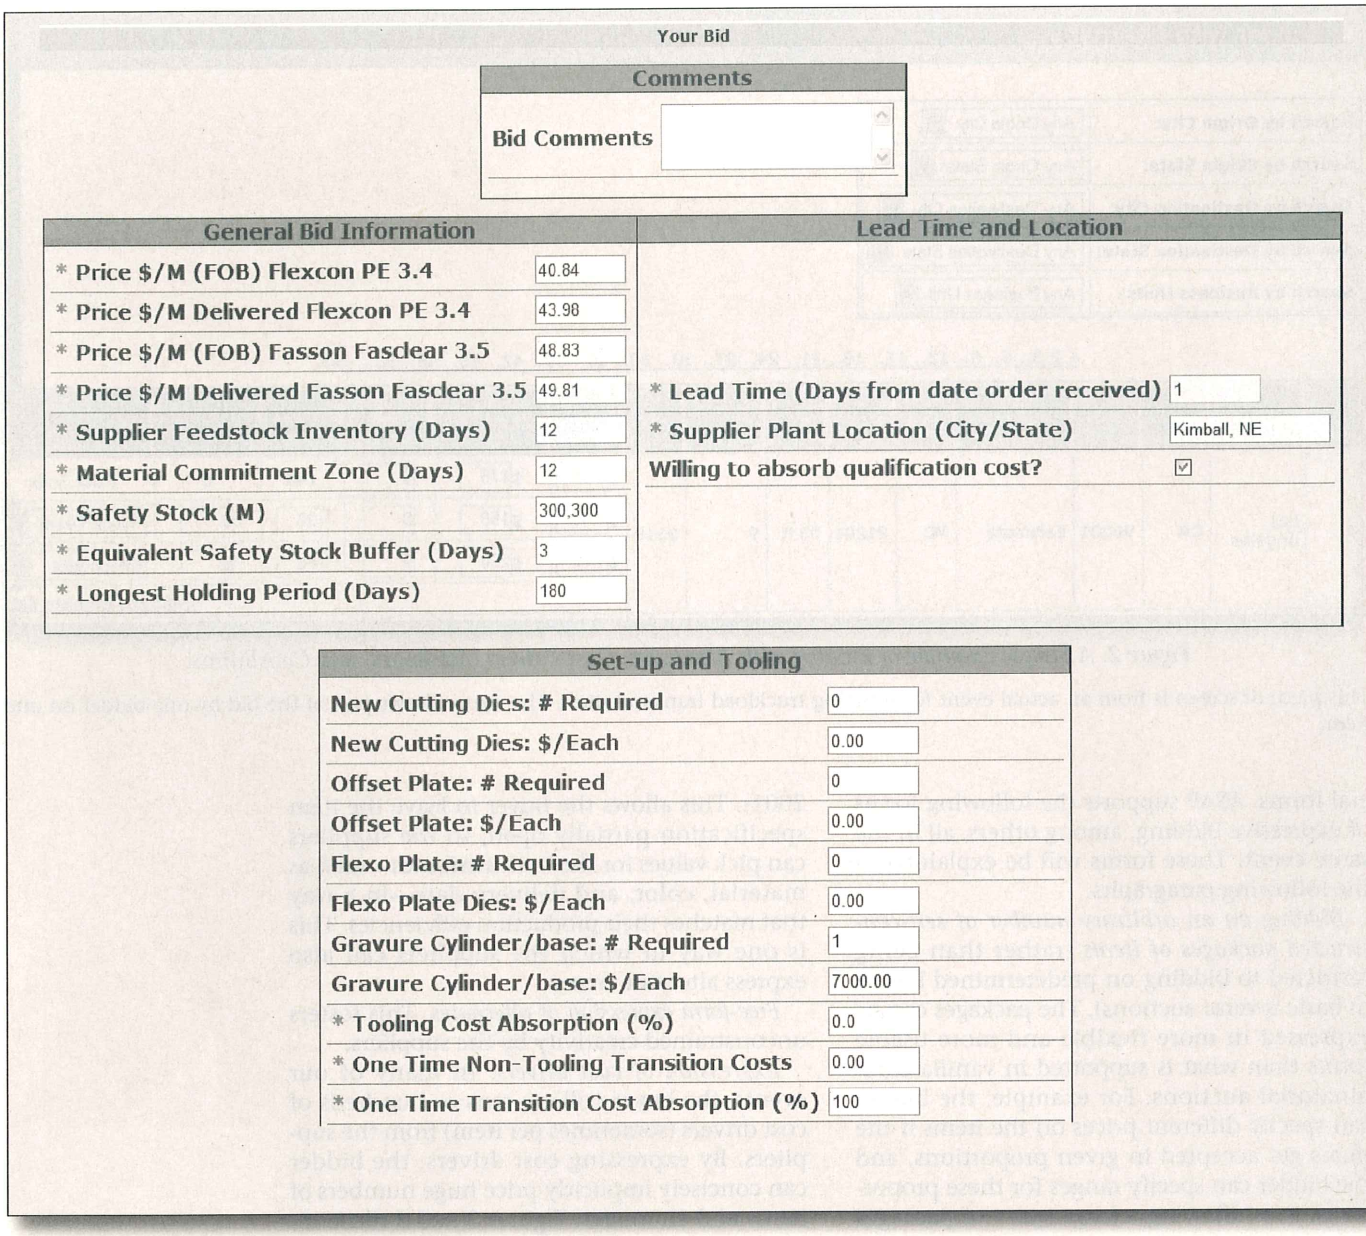

Figure 3. An Example of Bidding with Cost Structures and Attributes.

These screens are part of an actual event where printed labels were sourced.

Benefits of Expressive Bidding. The main benefit of expressive bidding is that it leads to a Pareto improvement in the allocation. In business terms, it creates a win-win between the buyer and the suppliers. There are several reasons for this.

First, because the suppliers and the buyer can express their preferences completely (and easily), the market mechanism can make better (economically more efficient and less wasteful) allocation decisions, which translates to higher societal welfare. In other words, the method yields better matching of supply and demand because they are expressed in more detail. The savings do not come from lowering supplier margins, but from reducing economic inefficiency. With expressive bidding, the suppliers can offer specifically what they are good at, and at lower prices because they end up supplying in a way that is economical for them. (They can consider factors such as production costs and capacities, raw material inventories, market conditions, competitive pressures, and strategic initiatives.) This creates a win-win solution between the suppliers and the buyer. For example, in the sourcing of transportation services, a substantial increase in economic efficiency comes from bundling multiple deliveries in one route (back-haul deliveries and multileg routes). This reduces empty driving, leading to lower transportation costs and yielding environmental benefits as well: lower fuel consumption, less driver time, less frequent need to replace equipment, and less pollution.

Second, suppliers avoid exposure risks. In traditional inexpressive markets, the suppliers face exposure problems when bidding. That makes bidding difficult. To illustrate this point, consider a simple auction of two trucking tasks: the first from Pittsburgh to Los Angeles, and the second from Los Angeles to Pittsburgh. If a carrier wins one of the tasks, the carrier has to factor in the cost of driving the other direction empty. Say that the cost for the task then is \$1.60 per mile. On the other hand, if the carrier gets both tasks, the carrier does not have to drive empty, and the cost is \$1.20 per mile. When bidding for the first task in an inexpressive auction, it is impossible to say where in the \$1.20–\$1.60 range the carrier should bid, because the cost for the first task depends on whether the carrier gets the second task, which in turn depends on how other carriers will bid. Any bid below \$1.60 exposes the carrier to a loss in case the carrier cannot profitably win the second task. Similarly, bidding above \$1.20 may cause the carrier to lose the deal on the first task although it would be profitable to take on that task in case the carrier wins the second task. In an expressive auction, the buyer can price each of the tasks separately, and price the package of them together, so there is no exposure problem. (For example, the buyer can bid \$1.60 per mile for the first task, \$1.60 per mile for the second task, and \$1.20 per mile for the package of both tasks. Of course, the buyer can also include a profit margin.) Therefore bidding is easier: the bidder does not have to speculate what other suppliers will bid in the later auctions. Also, the tasks get allocated optimally because no bidder gets stuck with an undesirable bundle, or misses the opportunity to win when the bidder is the most efficient supplier. Furthermore, when there is an exposure problem, the suppliers hedge against it by higher prices. Removal of the bidders' exposure problems thus also lowers the buyer's procurement cost.

Third, by expressive bidding with side constraints (such as capacity constraints), each supplier can bid on all bundles of interest without being exposed to winning so much that handling the business will be unprofitable or even infeasible. This again makes bidding easier because—unlike in inexpressive markets—the supplier does not have to guess which packages to commit his capacity to. (In an inexpressive market, making that guess requires counterspeculating what the other suppliers are going to bid, because that determines the prices at which this supplier can win different alternative packages.) This also leads to Pareto improvements in the allocation compared to inexpressive markets because in those markets each bidder needs to make guesses as to what parts of the business the bidder should bid on, and those parts might not be the parts for which the bidder really is the most efficient supplier.

Fourth, expressive bidding allows more straightforward participation in markets because the strategic counterspeculation issues that are prevalent in noncombinatorial markets are removed, as discussed above. This leads to wider access of the benefits of e-commerce because less experienced market participants are raised to an equal playing field with experts. This yields an increase in the number of market participants, which itself leads to further economic efficiency and savings in sourcing costs. Broader access also stems from the buyer not lotting the items and thus facilitating competition from small suppliers as well.

Fifth, in basic reverse auctions, the buyer has to prebundle items into lots, but the buyer cannot construct the optimal lotting because it depends on the suppliers' preferences. With expressive commerce, items do not have to be prebundled. Instead, the market determines the optimal lotting (specifically, the optimizer determines the optimal allocation based on the expressive bids from the suppliers and the expressions of constraints and preferences from the buyer). This way, the economically most efficient bundling is reached, weeks are not wasted on prebundling, and suppliers that are interested in different bundles compete. As a side effect, small suppliers' bids together end up competing with large suppliers.

Sixth, expressive bidding fosters creativity and innovation by the suppliers. This aspect is highly prized by both the suppliers and buyers. It can also yield drastic savings for the buyer due to creative construction of lower-cost alternates.

Overall, expressive bidding yields both lower prices and better supplier relationships. In addition to the buyers (our customers), also suppliers are providing very positive feedback on the approach. They especially like (1) that they also benefit (unlike in traditional reverse auctions where their profit margins get squeezed), (2) that they can express their production efficiencies, and (3) that they can express differentiation and creative offers. In fact, suppliers like expressive commerce so much that they agree to participate in expressive commerce even in events that they boycotted when basic reverse auctions had been attempted. Furthermore, perhaps the best indication of supplier satisfaction is the fact the suppliers are recommending the use of CombineNet to buyers.

The benefits of expressiveness can be further enhanced by multiple buyers conducting their sourcing in the same event. This provides an opportunity for the bidders to bundle across the demands of the buyers and also mitigates the exposure risks inherent in participating in separate events. As an example, in Spring 2005 CombineNet conducted an event where Procter & Gamble and its two largest customers, Walmart and Target, jointly sourced their North America–wide truckload transportation services for the following year (Sandholm et al. 2006). This enabled the carriers to construct beneficial backhaul deliveries and multileg routes by packaging trucking lanes across the demand of the three buyers. (This was a huge event. Procter & Gamble's volume alone exceeded $885 million.)

Expressive Allocation Evaluation

The second half of expressive commerce is expressive allocation evaluation, where the buyer expresses preferences over allocations using a rich, precise, and compact language that is natural in the buyer's business. It can be used to express legal constraints, business rules, prior contractual obligations, and strategic considerations.

In our experience, different types of side constraints are a powerful form of expressiveness for this purpose. For example, the buyer can state: "I don't want more than 200 winners (in order to avoid overhead costs)," "I don't want any one supplier to win more than 15 percent (in order to keep the supply chain competitive for the long term)," "I want minority suppliers to win at least 10 percent (because that is the law)," "Carrier X has to win at least $3 million (because I have already agreed to that)," and so on. ASAP supports hundreds of types of side constraints.

ASAP also has a rich language for the buyer to express how item attributes (such as delivery date or transshipment specifications) and supplier attributes (such as reputation) are to be taken into account when determining the allocation of business (Sandholm and Suri 2006).

A professional buyer—with potentially no background in optimization—can set up a scenario in ASAP by adding constraints and preferences through an easy-to-use web-based interface. (A simple example is shown in Figure 4.) To set up each such expression, the buyer first chooses the template expression (for example, "I don't want more than a certain number of winners," or "I want to favor incumbent suppliers by some amount") from a set of expressions that have been deemed potentially important for the event in question. The buyer then selects the scope to which that expression should apply: everywhere, or to a limited set of items, bid rounds, product groups, products, sites, and business groups. Finally, the buyer selects the exact parameters of the constraint, for example, exactly how many winners are allowed. Constraints and preferences can also be uploaded from business rule databases. Once the buyer has defined the scenario consisting of side constraints and preferences, the buyer calls the optimizer in ASAP to find the best allocation of business to suppliers under that scenario.

ASAP takes these high-level supply and demand expressions, automatically converts them into an optimization model, and uses sophisticated tree search algorithms to solve the model. CombineNet has faced scenarios with more than 2.6 million bids (on 160,000 items, multiple units of each) and more than 300,000 side constraints, and solved them to optimality.

Benefits of Expressive Allocation Evaluation. Through side constraints and preference expressions, the buyer can include business rules, legal constraints, logistical constraints, and other operational considerations to be taken into account when determining the allocation. This makes the auction's allocation implementable in the real world: the plan and execution are aligned because the execution considerations are captured in the planning.

Second, the buyer can include prior (for example, manually negotiated) contractual commitments into the optimization. This begets a sound hybrid between manual and electronic negotiation. For example, the buyer may have the obligation that a certain supplier has to be allocated at least 80 truckloads. The buyer can specify this as a side constraint in ASAP, and ASAP will decide which 80 truckloads (or more) are the best ones to allocate to that supplier in light of all other offers, side constraints, and preferences. This again makes the allocation implementable. (A poor person's way of accomplishing that would be to manually earmark some of the business to the prior contracts. Naturally, allowing the system to do that earmarking with all the pertinent information in hand yields better allocations.)

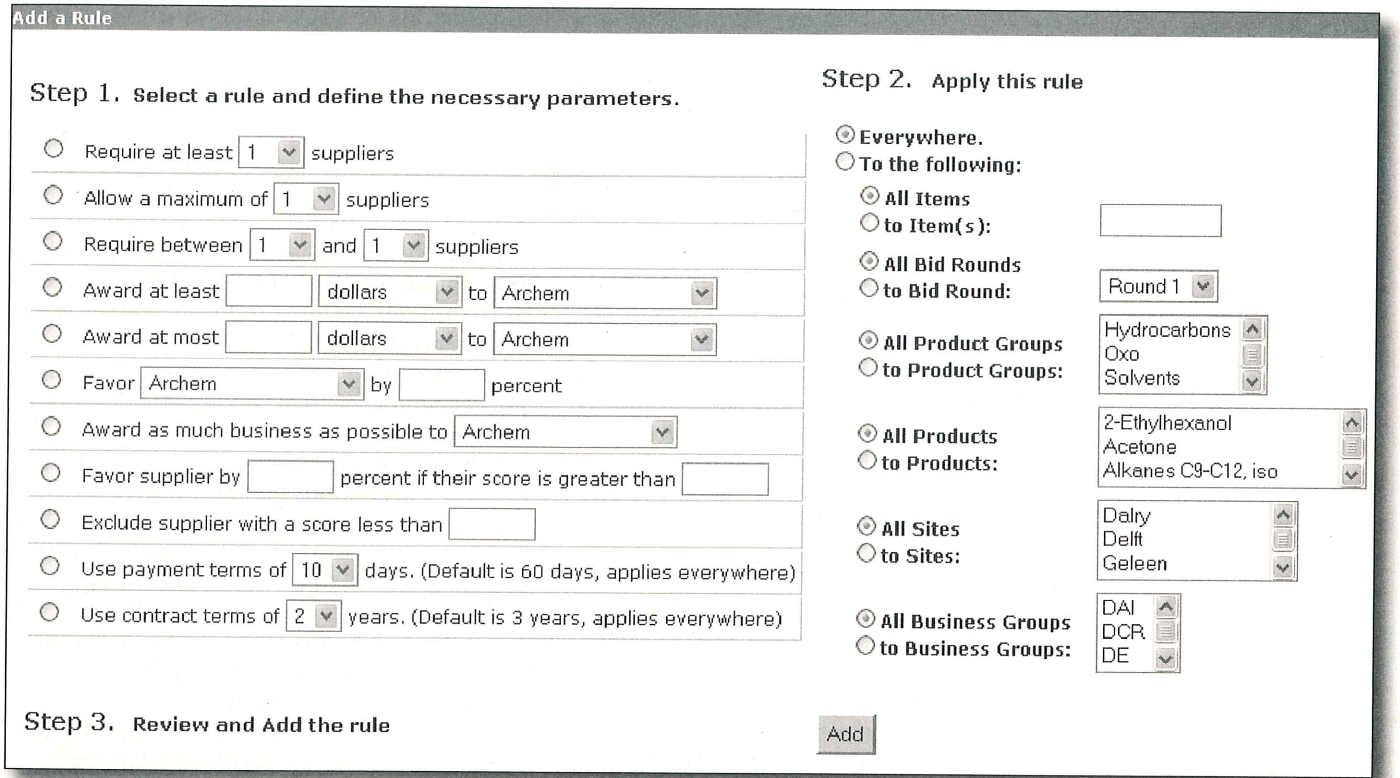

Figure 4. A User Interface for Expressive Allocation Evaluation by the Bid Taker.

Every sourcing event has different expressiveness forms (selected from a library or preconfigured in a template). This particular sourcing event was one of the simpler ones in that relatively few expressions of constraints and preferences were included in the user interface. The buyer uses the interface as follows. First, on the left (Step 1), the buyer selects which one of the expressiveness forms the buyer wants to work on and sets the parameters for that form. Then, on the right (Step 2) the buyer selects the scope to which this rule is to be applied. In Step 3, the buyer presses the "Add" button to add the rule, and a restatement of the rule appears in natural language for verification (not shown). The buyer can then add more rules to the same scenario by repeating this process. Finally, the buyer presses the "Optimize" button (not shown) to find the optimal allocation for the scenario. This triggers the automated formulation of all these constraints and preferences—together with all the bid information that the bidders have submitted—into an optimization problem, and the solving of the problem through advanced tree search.

Third, the buyer obtains a quantitative understanding of the trade-offs in his supply chain by conducting what-if *scenario navigation;* that is, by changing side constraints and preferences and reoptimizing, the buyer can explore the trade-offs in an objective manner. For example, the buyer may add the side constraint that the supply base be rationalized from 200 to 190. The resulting increase in procurement cost then gives the buyer an understanding of the trade-off between cost and practical implementability. As another example, the buyer might ask: If I wanted my average supplier delivery-on-time rating to increase to 99 percent, how much would that cost? As a third example, the buyer might see what would happen if the buyer allowed a supplier to win up to 20 percent of the business instead of only 15 percent. The system will tell the buyer how much the procurement cost would decrease. The buyer can then decide whether the savings outweighs the added long-term strategic risks such as vulnerability to that supplier's default and the long-term financial downside of allowing one supplier to become dominant. In ASAP, the buyer can change/add/delete any number of side constraints and preferences in between optimizations.

Fourth, quantitative understanding of the trade-offs also fosters stakeholder alignment on the procurement team, because the team members with different preferences can discuss based on facts rather than opinions, philosophies, and guesswork.

Time to Contract in Expressive Commerce

The time to contract is reduced from several months to weeks because no manual lotting is required, all suppliers can submit their offers in parallel, what-if scenarios can be rapidly generated and analyzed, and the allocation is implementable as is. This causes the cost savings to start to accrue earlier, and decreases the human hours invested.

Automated Item Configuration

An additional interesting aspect of bidding with cost drivers and alternates (for example, using attributes) is that the market-clearing (aka winner determination) algorithm not only decides who wins but also ends up optimizing the configuration (setting of attributes) for each item. In deciding this, the optimizer, of course, also takes into account the buyer's constraints and preferences.

Automated Supply Chain Configuration

In many sourcing events, the winner determination also ends up optimizing the supply chain multiple levels upstream from the buyer. For example, in an event where Procter & Gamble sourced in-store displays using our hosting service and technology, we sourced items from different levels of the supply chain in one event: buying colorants and cardboard of different types, buying the service of printing, buying the transportation, buying the installation service, and so on (Sandholm et al. 2006). Some suppliers made offers for some of those individual items while others offered complete ready-made displays (which are, in effect, packages of the lower-level items), and some bid for partial combinations. The market clearing determined the lowest-cost (adjusted for the Procter & Gamble's constraints and preferences) solution and thus, in effect, configured the supply chain multiple levels upstream.

Expressive Commerce as a Generalization of Combinatorial Auctions

A relatively simple early form of expressive commerce was a *combinatorial reverse auction* (Sandholm et al. 2002), where the only form of expressiveness that the suppliers have is package bidding, and the buyer has no expressiveness. A predecessor of that was a *combinatorial auction* where the bidders are the buyers (and there is only one unit of each item and no side constraints). Combinatorial auctions (Rassenti, Smith, and Bulfin 1982; Sandholm 1993; Sandholm 2002b; Ledyard, Porter, and Rangel 1997; Rothkopf, Pekec, and Harstad 1998; Kwasnica et al. 2005; Sandholm et al. 2005; Sandholm and Suri 2003; Hoos and Boutilier 2000; Boutilier 2002; and deVries and Vohra 2003) enable bidders to express complementarity among items (the value of a package being more than the sum of its parts) through package bids. Substitutability (the value of a package being less than the sum of its parts) can also be expressed in some combinatorial auctions, usually using different languages for specifying mutual exclusivity between bids (Sandholm 2002a; Fujishima, Leyton-Brown, and Shoham 1999; Sandholm 2002b; Nisan 2000; and Hoos and Boutilier 2001).

Expressiveness leads to more economical allocations of the items because bidders do not get stuck with partial bundles that are of low value to them. This has been demonstrated, for example, in auctions for bandwidth (McMillan 1994, McAfee and McMillan 1996), transportation services (Sandholm 1993, Sandholm 2000, Sandholm 1991, and Caplice and Sheffi 2003), pollution rights, airport landing slots (Rassenti, Smith, and Bulfin 1982), and carrier-of-last-resort responsibilities for universal services (Kelly and Steinberg 2000).

However, package bids and exclusivity constraints are too impoverished a language for real-world sourcing. While any mapping from bundles to real numbers can be expressed in that language in principle, the real-world preferences in sourcing cannot be easily, naturally, and concisely expressed in it. Starting in 1997, we tackled this challenge and generalized the approach to expressive commerce, with the language constructs discussed above. Similar approaches have recently been adopted by others but only for drastically less complex (orders of magnitude smaller and less expressive) events (Hohner et al. 2003, Metty et al. 2005).

The use of our richer expressiveness forms (rather than mere canonical package bids with exclusivity constraints) is of key importance for several reasons: (1) Bidders can express their preference in the language that is natural in their domain. (2) Bidders can express their preferences concisely. To illustrate this point, consider the following simple example. A bidder has no production efficiencies and thus has a price for each item regardless of what other items the buyer produces. However, the bidder has a capacity constraint. In our bidding language, the bidder can simply express a price for each item and a capacity constraint. In contrast, in the classical combinatorial auction bidding languages, the supplier would have to submit bids for an exponential number of packages. (3) Due to the conciseness, the bids are easy to communicate to the bid taker. (4) Our bidding constructs maintain the natural structure of the problem (rather than normalizing the structure away into a format that only allows package bids with exclusivity constraints). The clearing algorithms take advantage of that structure in many ways, for example, in generating cutting planes, deciding what variables to branch on, and so on.

Tree Search to Enable Expressive Commerce

A significant challenge in making expressive commerce a reality is that the expressiveness makes the problem of allocating the business across the suppliers an extremely complex combinatorial optimization problem. Specifically, the clearing problem (aka winner determination problem) is that of deciding which bids to accept and reject (and to what extent in the case of partially acceptable bids) so as to minimize sourcing cost (adjusted for preferences) subject to satisfying all the demand and all side constraints. Even in the vanilla combinatorial reverse auction where the only form of bidding is package bidding, and no side constraints or preferences are allowed, the clearing problem is NP-complete and inapproximable in the worst case in polynomial time (Sandholm et al. 2002). Expressive commerce is a much richer problem; thus the NP-hardness and inapproximability carry over. (Müller, Lehmann, and Sandholm [2006] review the worst-case complexity of the clearing problem of different variants of combinatorial auctions.) Thus sophisticated techniques are required.

In fact, prior to ASAP, no technology was capable of solving clearing problems of the scale and expressiveness that our customers wanted to be able to support; for example, Hohner et al. (2003) found integer programming techniques to be effective for problems only as large as 500 items and 5000 bids. In 2001, P&G gave us a trial instance of trucking services sourcing that took a competing optimization product 30 minutes to solve. ASAP solved it optimally in 9 seconds. While that was already a decisive speed difference, since that time our technology development has yielded a further speed improvement of two to three orders of magnitude.

There is significant structure in the expressive commerce problem instances, and it is paramount that the optimizer be able to take advantage of the structure. Mixed integer programming (MIP) techniques for tree search are quite good at this, and ASAP takes advantage of them. However, the techniques embodied in the leading general-purpose MIP solvers alone are not sufficient for the clearing problem.

ASAP uses sophisticated tree search to find the optimal allocation. Given that the problem is NP-complete, in the worst case the run time is superpolynomial in the size of the input (unless P = NP). However, in real-world sourcing optimization, the algorithms run extremely fast: the median run time is less than a second, and the average is 20 seconds, with some instances taking days. The algorithms are also anytime algorithms: they provide better and better solutions during the search process.

I began the algorithm development in 1997, and CombineNet now has 16 people working on the algorithms, half of them full time. The team has tested hundreds of techniques (some from the AI and operations research literature and some invented at CombineNet) to see which ones enhance speed on expressive commerce clearing problems. Some of the techniques are specific to market clearing, while others apply to combinatorial optimization more broadly. We published the first generations of our search algorithms (Sandholm 2002a, Sandholm and Suri 2003, and Sandholm et al. 2005). The new ideas in these algorithms included different formulations of the basic combinatorial auction clearing problem (branching on items [Sandholm 2002a], branching on bids [Sandholm and Suri 2003, Sandholm et al. 2005], and multivariable branching [Gilpin and Sandholm 2007]), upper and lower bounding across components in dynamically detected decompositions (Sandholm and Suri 2003, Sandholm et al. 2005), sophisticated strategies for branch question selection (Sandholm 2006, Sandholm 2002a, Sandholm and Suri 2003, and Sandholm et al. 2005), dynamically selecting the branch selection strategy at each search node (Sandholm 2006, Sandholm et al. 2005), the information-theoretic approach to branching in search (Gilpin and Sandholm 2007), sophisticated lookahead techniques (Sandholm 2006, Gilpin and Sandholm 2007), solution seeding (Sandholm 2006), primal heuristics (Sandholm 2006, Sandholm et al. 2005), identifying and solving tractable cases at nodes (Sandholm and Suri 2003; Sandholm et al. 2005; Sandholm 2006; and Conitzer, Derryberry, and Sandholm 2004), techniques for exploiting part of the remaining problem falling into a tractable class (Sandholm 2006, Sandholm and Suri 2003), domain-specific preprocessing techniques (Sandholm 2002a), fast data structures (Sandholm 2002a, Sandholm and Suri 2003, and Sandholm et al. 2005), methods for handling reserve prices (Sandholm 2002a, Sandholm and Suri 2003), and incremental winner determination and quote computation techniques

(Sandholm 2002a). Sandholm (2006) provides an overview of the techniques.

We have also invented a host of techniques in the search algorithms that we have decided to keep proprietary for now. They include different formulations of the clearing problem, new branching strategies, custom cutting plane families, cutting plane generation and selection techniques, and so on.

Since 2002 we have also been using machine learning methods to predict how well different techniques will perform on the instance at hand. (For this purpose, the instance is represented by about 50 hand-selected numeric features.) This information can be used to dynamically select the technique for the instance at hand, to give time estimates to the user, and so on. Our solver has several dozen important parameters and each of them can take on several values. Therefore, our machine learning approach of setting the parameters well on an instance by instance basis is significantly more challenging and more powerful than using machine learning to select among a handful of hardwired solvers, an approach that has been successfully pursued in academia in parallel (for example, Leyton-Brown, Nudelman, and Shoham [2006]).

While the literature on combinatorial auctions has mainly focused on a variant where the only form of expressiveness is package bidding (sometimes supplemented with mutual exclusion constraints between bids), in our experience with sourcing problems the complexity is dominated by rich side constraints. Thus we have invested significant effort into developing techniques that deal with side constraints efficiently. CombineNet has faced several hundred different types of real-world side constraints. ASAP supports all of them. We abstracted them into eight classes from an optimization perspective so the speed improvements that we build into the solver for a type of side constraint get leveraged across all side constraint types within the class.

The resulting optimal search algorithms are often 10,000 times faster than those of others. The main reason is that CombineNet specializes on a subclass of MIP problems and has 32,000 real-world instances on which to improve its algorithms. The speed has allowed our customers to handle drastically larger and more expressive sourcing events. The events have sometimes had more than 2.6 million bids (on 160,000 items, multiple units of each) and more than 300,000 side constraints.

The state-of-the-art general-purpose MIP solvers are inadequate, as is, also due to numeric instability. They err on feasibility, optimality, or both, on about 4 percent of the sourcing instances. We have invested significant effort on stability, yielding techniques that are significantly more robust.

Hosted Optimization for Sourcing Professionals

CombineNet's backend clearing engine, ClearBox, is industry independent, and the interface to it is through our Combinatorial Exchange Description Language (CEDL), an XML-based language that allows ClearBox to be applied to a wide variety of applications by CombineNet and its partners (see figure 5).

Intuitive web-based interfaces designed for the buyer and for the suppliers bring the power of optimization to users with expertise in sourcing, not in optimization. The users express their preferences through interfaces that use sourcing terminology. The interfaces support simple click-through interaction rather than requiring the user to know any syntax. The approach allows sourcing professionals to do what they are best at (incorporating sourcing knowledge such as strategic and operational considerations) and the optimizer to do what it is best at (sifting through huge numbers of allocations to pick the best one).

For every event, separate front-ends are configured that support only those bidding and allocation evaluation features that are appropriate for that event. This makes the user interfaces easier and more natural to use by sourcing professionals. User training typically takes a few hours. New front ends typically take from a few hours to a couple of weeks to go from project specification to deployment.

The user interfaces feed CEDL into ClearBox, and ClearBox then automatically formulates the optimization problem for the search algorithms. This contrasts with the traditional mode of using optimization, where a consultant with optimization expertise builds the model. The automated approach is drastically faster (seconds rather than months) and avoids errors.

Our web-based products and application service provider (ASP) business model make optimization available on demand. No client-side software installation is necessary. This also avoids hardware investments by customers. We buy the hardware and leverage it across customers, each with temporary load. (On many instances the search trees exceed 2 gigabytes of RAM, rendering 32-bit architectures unusable and requiring a 64-bit architecture.) See figure 6. The ASP model allows us quickly and transparently to tune our algorithms and to provide

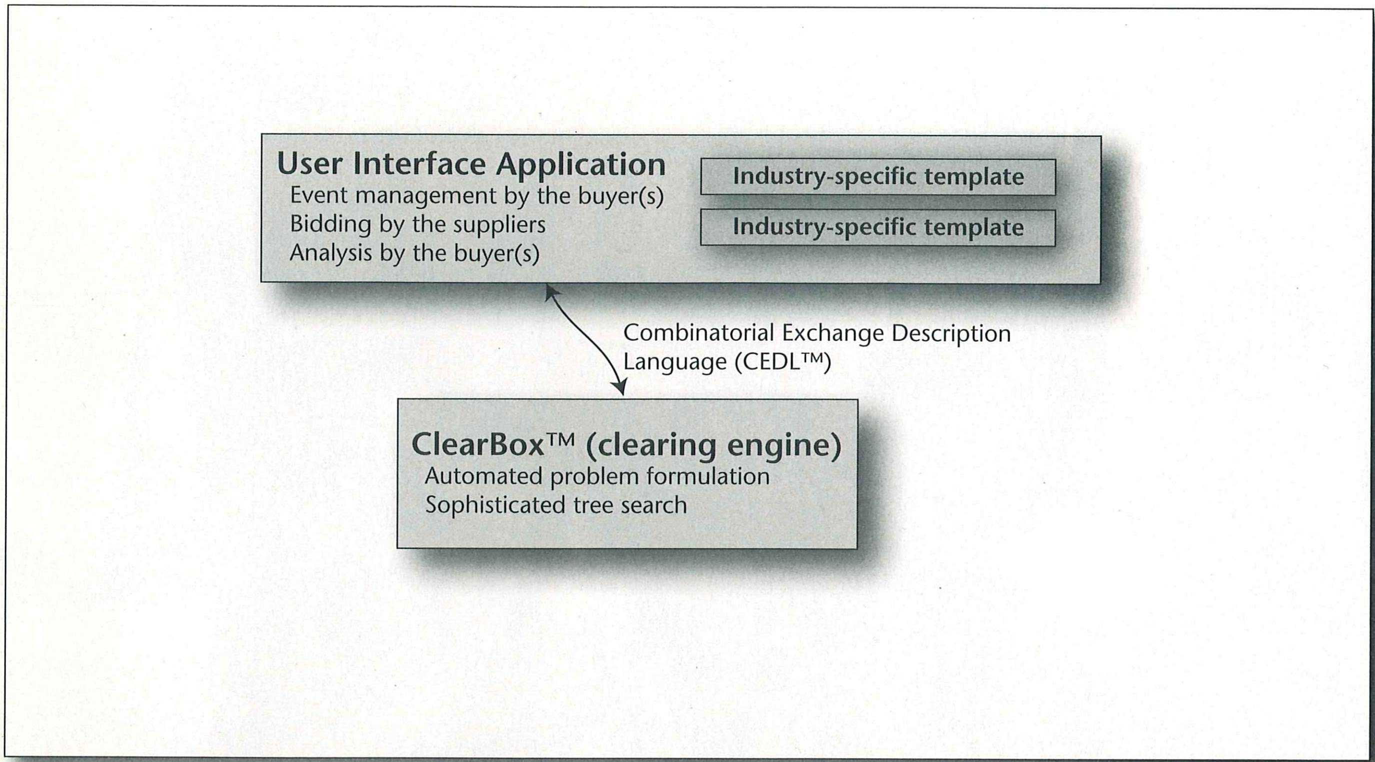

Figure 5. Advanced Sourcing Application Platform (ASAP).

The platform is hosted on a server farm with multiple instantiations of each component. ASAP also includes modules for clearing management, server farm management, secure databases, and so on (not shown).

enhancements to all customers simultaneously. We also offer the technology through consulting firms.

Market Design

While we attribute the bulk of the savings to the application of optimization to sourcing, another important factor is market design: what forms of expressiveness are allowed, what form of feedback is given to bidders during the event, and so on. ASAP supports sealed bid events (winners are determined at the end), events that have a (usually small) number of rounds (winners are determined and feedback provided at the end of each round), and "live" events (winners are determined and feedback provided every time any participant expresses anything new).

Scenario Navigation

The buyer is typically not an individual but an organization of several individuals with different preferences over allocations. Finance people want low sourcing cost, plant managers want small numbers of suppliers, marketing people want a high average carrier-delivery-on-time rating, and so on. ASAP enables the organization to better understand the available trade-offs. Once the bids have been collected, the buyer conducts scenario navigation. At each step of that process, the buyer specifies a set of side constraints and preferences (these define the scenario), and runs the optimizer to find an optimal allocation for that scenario. This way the buyer obtains a quantitative understanding of how different side constraints and preferences affect the sourcing cost and all other aspects of the allocation.

CombineNet has found that a buying organization will navigate an average of 100 scenarios per sourcing event. (The maximum seen to date had 1107.) To navigate such large numbers of scenarios, fast clearing is paramount.

Rapid clearing allows scenario navigation to be driven by the actual data (offers). In contrast, most prior approaches required the scenario (side constraints and preferences, if any) to be defined prior to analysis; there were insufficient time and expert modeling resources to try even a small number of alternative scenarios. Data-driven approaches are clearly superior because the actual offers provide accurate costs for the various alternative scenarios.

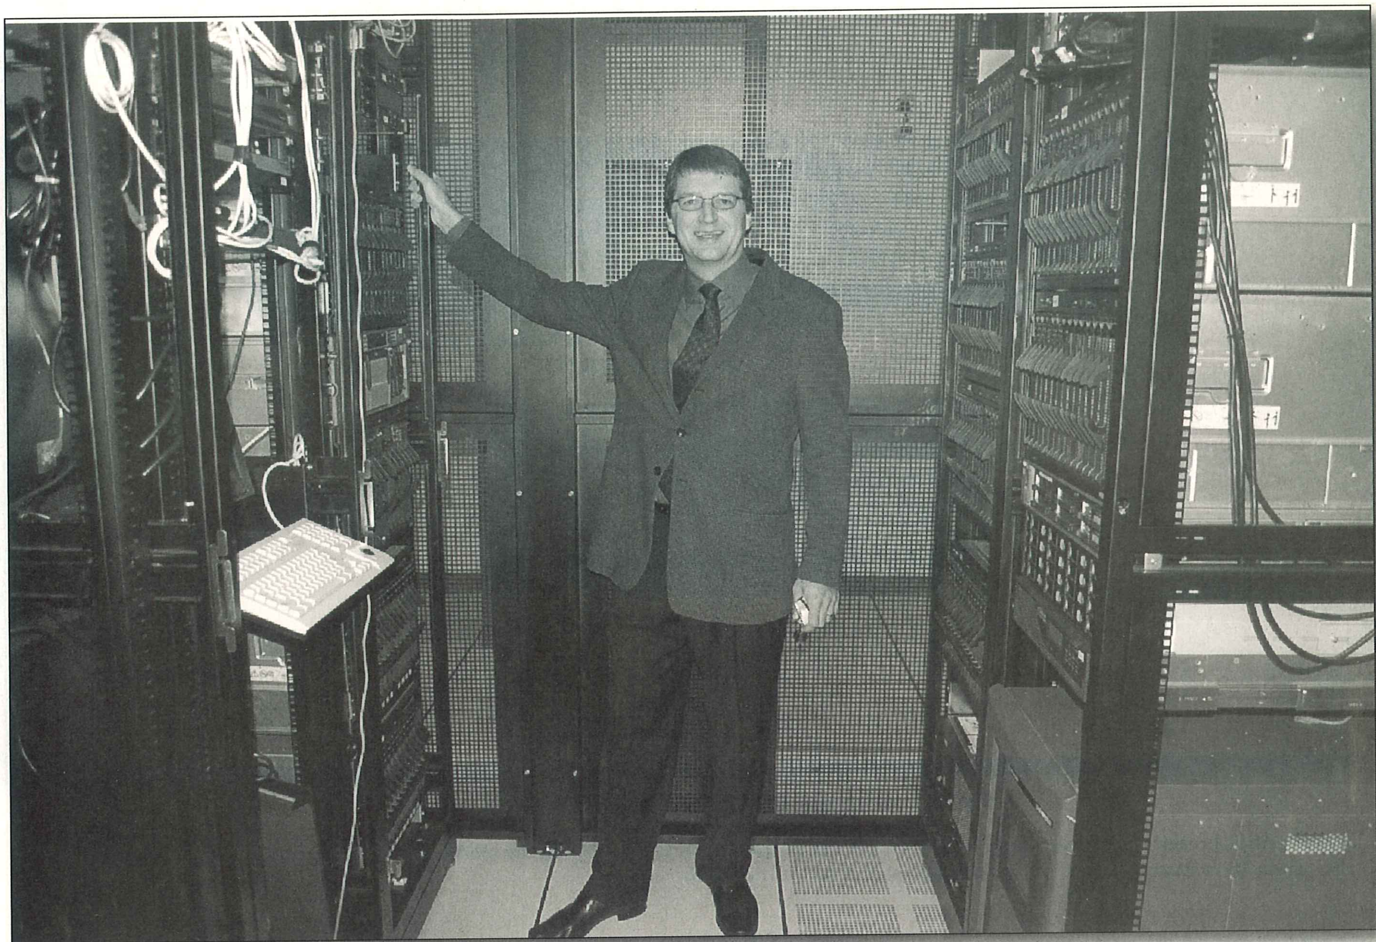

Figure 6. CombineNet's Hosting Architecture Now Includes Hundreds of Computers, Hardware Load Balancers, Redundant Internet, and So On.

Other aspects of the hosting include 24/7 monitoring of the hosting infrastructure, backup power from a generator, measures against overheating and fire, and physical security against unauthorized entry.

The next generation of ASAP will also support *automated scenario navigation.* Compared to basic scenario navigation, discussed above, it enables a more systematic and less wasteful navigation of the scenario space. The system queries the sourcing team about their preferences, using, for example, trade-off queries ("how much hassle would an extra supplier be in dollars?—give me an upper or lower bound") and comparison queries ("which of these two allocations do you prefer?"). The system decides the queries to pose in a data-directed way so as to ask the team to refine its preferences only on an as-needed basis. (This is desirable because internal negotiation in the team is costly in terms of time and goodwill.) Specifically, based on all the offers that the suppliers have submitted, and all answers to previous queries, the system strives to minimize maximum regret. At each iteration of automated scenario navigation, the system finds a robust solution that minimizes maximum regret (the regret is due to the fact that the sourcing team has not fully specified its preferences, so for some preferences that are still consistent with the answers so far, the system's recommended allocation is not optimal). As the other step of each iteration, the system poses a query to refine the team's preferences in order to be able to reduce the maximum regret further. The maximum regret also provides a quantitative measure of when further negotiation within the team is no longer worth it, and the team should implement the current robust allocation. CombineNet pioneered automated scenario navigation, including its different design dimensions and algorithms (Boutilier, Sandholm, and Shields 2004). The optimization problem of finding the most robust allocation is even more complex than the clearing problem discussed in most of this paper. We have developed a prototype of automated scenario navigation, and will solicit customer feedback soon.

Impact

The new sourcing paradigm and technology has already had significant impact. The technology development began in 1997 and CombineNet was founded in 2000. Between December 2001 and now (December 2006), CombineNet has used ASAP to host 447 highly combinatorial procurement events, totaling a spend of $35 billion. The 60 plus buyer companies were mostly among the Global 1000. A total of more than 12,000 supplier companies bid in our system. The individual events ranged from $2 million to $1.6 billion, representing the most complex combinatorial auctions ever conducted. They spanned a broad range of categories such as transportation (truckload, less-than-truckload, ocean freight, dray, bulk, intermodal, small parcel, air freight, train, fleet, freight forwarding, and other), direct materials (sugars and sweeteners, meat, vegetables, honey, starches, colorants, fibers and nonwovens, steel, fasteners, solvents, chemicals, casings, resins, and polymers), packaging (cans and ends, corrugates, corrugated displays, flexible film, folding cartons, labels, foam trays and pads, caps and closures, shrink and stretch films, bags, pulp, pallets, and printed instructions), indirect materials (management, repair, and operations, [MRO] [electrical supplies, filters, pipes and valves and fittings, power transmissions, pumps, safety supplies, office supplies, lab supplies, file folders, solvents, and furnishings], chemicals [cylinder gasses, fuels, and other], technology [laptops and desktops and cameras], leased equipment, fleet vehicles, and promotional items, and services (security, janitorial, legal, patent and trademark, consulting, equipment maintenance, temp labor, marketing, customization, insurance, shuttling and towing, warehousing, prepress, and advertising), and healthcare (pharmaceuticals as well as medical and surgical equipment and supplies), and telecommunication (sourcing wireless plans for employees of companies).

On the $35 billion spend, CombineNet delivered hard-dollar savings of $4.4 billion to its customers in lowered sourcing costs. The savings were measured compared to the prices that the buyer paid for the same items the previous time the buyer sourced them (usually 12 months earlier). The $4.4 billion is the implementable savings that the system yielded after the buyer had applied side constraints and preferences; the unconstrained savings (that is, the savings arising from expressive bidding before the buyer expresses constraints and preferences) was $5.4 billion. The savings figure is remarkable especially taking into account that during the same time period, the prices in the largest segment, transportation, increased by 6–9 percent in the market overall.

The savings number does not include the savings obtained by suppliers, which are harder to measure because the suppliers' true cost structures are proprietary. However, there is strong evidence that the suppliers also benefited, so a win-win was indeed achieved. First, suppliers that have participated in expressive commerce events are recommending the use of that approach to other buyers. Second, on numerous occasions, suppliers that boycotted reverse auctions came back to the "negotiation table" once expressive commerce was introduced. Third, suppliers are giving very positive feedback about their ability to express differentiation and provide creative alternatives.

The savings number also does not include savings that stem from reduced effort and compression of the event timeline from months to weeks or even days.

The cost savings were achieved while at the same time achieving the other advantages of expressive commerce discussed above, such as better supplier relationships (and better participation in the events), redesign of the supply chain, implementable solutions that satisfy operational considerations, and solutions that strike the trade-offs in a data-driven way and align the stakeholders in the buying organization. See also Sandholm et al. (2006) and case studies at www.CombineNet.com.

Today CombineNet employs 130 full time (about half of them in engineering) and a dozen academics as advisors. The company has operations on four continents. It is headquartered in Pittsburgh, Pennsylvania, USA, with European headquarters in Berlin, Germany. It also has offices in Brussels, Belgium; Hamburg, Germany; Tokyo, Japan; and Beijing, China.

Acknowledgments

I thank the CombineNet employees for helping make this vision a reality. Special thanks go to Subhash Suri, David Levine, Tony Bonidy, Paul Martyn, Andrew Gilpin, Rob Shields, Bryan Bailey, Andrew Fuqua, David Parkes, George Nemhauser, Craig Boutilier, and Egon Balas. CombineNet, Expressive Bidding, Expressive Competition, Expressive Commerce, Expressive Bid Taking, ClearBox, Combinatorial Exchange Description Language, and CEDL are trademarks of CombineNet, Inc.

References

Boutilier, C. 2002. Solving Concisely Expressed Combinatorial Auction Problems. In *Proceedings of the Eighteenth National Conference on Artificial Intelligence.* Menlo Park, CA: AAAI Press.

Boutilier, C.; Sandholm, T.; and Shields, R. 2004. Eliciting Bid Taker Nonprice Preferences in (Combinatorial) Auctions. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence.* Menlo Park, CA: AAAI Press.

Caplice, C., and Sheffi, Y. 2003. Optimization-Based Procurement for Transportation Services. *Journal of Business Logistics* 24(2): 109–128.

Conitzer, V.; Derryberry, J.; and Sandholm, T. 2004. Combinatorial Auctions with Structured Item Graphs. In *Proceedings of the Nineteenth National Conference on Artificial Intelligence.* Menlo Park, CA: AAAI Press.

de Vries, S., and Vohra, R. 2003. Combinatorial Auctions: A Survey. *INFORMS Journal on Computing* 15(3): 284–309.

Fujishima, Y.; Leyton-Brown, K.; and Shoham, Y. 1999. Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches. In *Proceedings of the Sixteenth*

International Joint Conference on Artificial Intelligence. San Francisco: Morgan Kaufmann Publishers.

Gilpin, A., and Sandholm, T. 2007. Information-Theoretic Approaches to Branching in Search. In *Proceedings of the Twentieth International Joint Conference on Artificial Intelligence.* Menlo Park, CA: AAAI Press.

Hohner, G.; Rich, J.; Ng, E.; Reid, G.; Davenport, A.; Kalagnanam, J.; Lee, H.; and An, C. 2003. Combinatorial and Quantity-Discount Procurement Auctions Benefit Mars, Incorporated and Its Suppliers. *Interfaces* 33(1): 23–25.

Hoos, H., and Boutilier, C. 2000. Solving Combinatorial Auctions Using Stochastic Local Search. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence.* Menlo Park, CA: AAAI Press.

Hoos, H., and Boutilier, C. 2001. Bidding Languages for Combinatorial Auctions. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence.* San Francisco: Morgan Kaufmann Publishers.

Kelly, F., and Steinberg, R. 2000. A Combinatorial Auction with Multiple Winners for Universal Services. *Management Science* 46(4): 586–596.

Kothari, A.; Parkes, D.; and Suri, S. 2003. Approximately Strategyproof and Tractable Multi-Unit Auctions. In *Proceedings, Fourth ACM Conference on Electronic Commerce,* 166–175. New York: Association for Computing Machinery.

Kwasnica, A.; Ledyard, J.; Porter, D.; and DeMartini, C. 2005. A New and Improved Design for Multiobject Iterative Auctions. *Management Science* 51(3): 419–434.

Ledyard, J.; Porter, D.; and Rangel, A. 1997. Experiments Testing Multiobject Allocation Mechanisms. *Journal of Economics and Management Strategy* 6(3): 639–675.

Leyton-Brown, K.; Nudelman, E.; and Shoham Y. 2006. Empirical Hardness Models for Combinatorial Auctions. Chapter 19 of *Combinatorial Auctions,* ed. P. Cramton, Y. Shoham, and R. Steinberg. Cambridge, MA: The MIT Press.

McAfee, P., and McMillan, J. 1996. Analyzing the Airwaves Auction. *Journal of Economic Perspectives* 10(1): 159–175.

McMillan, J. 1994. Selling Spectrum Rights. *Journal of Economic Perspectives* 8(3): 145–162.

Metty, T.; Harlan, R.; Samelson, Q.; Moore, T.; Morris, T.; Sorensen, R.; Schneur, A.; Raskina, O.; Schneur, R.; Kanner, J.; Potts, K.; and Robbins, J. 2005. Reinventing the Supplier Negotiation Process at Motorola. *Interfaces* 35(1): 7–23.

Müller, R.; Lehmann, D.; and Sandholm, T. 2006. The Winner Determination Problem. Chapter 12 of *Combinatorial Auctions,* ed. P. Cramton, Y. Shoham, and R. Steinberg. Cambridge, MA: The MIT Press.

Nisan, N. 2000. Bidding and Allocation in Combinatorial Auctions. In *Proceedings, First ACM Conference on Electronic Commerce,* 1–12. New York: Association for Computing Machinery.

Rassenti, S.; Smith, V.; and Bulfin, R. 1982. A Combinatorial Auction Mechanism for Airport Time Slot Allocation. *Bell Journal of Economics* 13(2): 402–417.

Rothkopf, M.; Pekec, A.; and Harstad, R. 1998. Computationally Manageable Combinatorial Auctions. *Management Science* 44(8): 131–1147.

Sandholm, T. 1991. A Strategy for Decreasing the Total Transportation Costs Among Area-Distributed Transportation Centers. In *Nordic Operations Analysis in Cooperation (NOAS-91): OR in Business.* Turku, Finland: Turku School of Economics.

Sandholm, T. 1993. An Implementation of fhe Contract Net Protocol Based on Marginal Cost Calculations. In *Proceedings of the Eleventh National Conference on Artificial Intelligence.* Menlo Park, CA: AAAI Press.

Sandholm, T. 2000. Issues in Computational Vickrey Auctions. *International Journal of Electronic Commerce* 4(3): 107–129.

Sandholm, T. 2002a. Algorithm for Optimal Winner Determination in Combinatorial Auctions. *Artificial Intelligence* 135(1–2): 1–54.

Sandholm, T. 2002b. eMediator: A Next Generation Electronic Commerce Server. *Computational Intelligence* 18(4): 656–676.

Sandholm, T. 2006. Optimal Winner Determination Algorithms. Chapter 14 of *Combinatorial Auctions,* ed. P. Cramton, Y. Shoham, and R. Steinberg. Cambridge, MA: The MIT Press.

Sandholm, T.; Levine, D.; Concordia, M.; Martyn, P.; Hughes, R.; Jacobs, J.; and Begg, D. 2006. Changing the Game in Strategic Sourcing at Procter & Gamble: Expressive Competition Enabled by Optimization. *Interfaces* 36(1), 55–68.

Sandholm, T., and Suri, S. 2001. Market Clearability. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence.* San Francisco: Morgan Kaufmann Publishers.

Sandholm, T., and Suri, S. 2006. Side Constraints and Nonprice Attributes in Markets. *Games and Economic Behavior* 55(2): 321–330.

Sandholm, T., and Suri, S. 2002. Optimal Clearing of Supply/Demand Curves. In *Proceedings of the Thirteenth Annual International Symposium on Algorithms and Computation (ISAAC),* Lecture Notes in Computer Science, Vol. 2518. Berlin: Springer.

Sandholm, T., and Suri, S. 2003. BOB: Improved Winner Determination in Combinatorial Auctions and Generalizations. *Artificial Intelligence* 145(1–2): 33–58.

Sandholm, T.; Suri, S.; Gilpin, A.; and Levine, D. 2002. Winner Determination in Combinatorial Auction Generalizations. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS).* New York: Association for Computing Machinery.

Sandholm, T.; Suri, S.; Gilpin, A.; and Levine, D. 2005. CABOB: A Fast Optimal Algorithm for Winner Determination in Combinatorial Auctions. *Management Science* 51(3): 374–390.

Smock, D. 2004. How P&G Buys Plastics. *Modern Plastics* (December 1).

Tuomas Sandholm is a professor in the Computer Science Department at Carnegie Mellon University and founder, chairman, and chief scientist of CombineNet, Inc. He received Ph.D. and M.S. degrees in computer science from the University of Massachusetts at Amherst in 1996 and 1994. He earned an M.S. (B.S. included) with distinction in industrial engineering and management science from the Helsinki University of Technology, Finland, in 1991. He has published more than 250 papers on electronic commerce, game theory, artificial intelligence, multiagent systems, auctions and exchanges, automated negotiation and contracting, coalition formation, voting, safe exchange, normative models of bounded rationality, resource-bounded reasoning, machine learning, networks, and combinatorial optimization. Several of his systems have been commercially fielded. He received the NSF CAREER Award in 1997, the inaugural ACM Autonomous Agents Research Award in 2001, the Alfred P. Sloan Foundation Fellowship in 2003, and the IJCAI Computers and Thought Award in 2003.