

Expressiveness in mechanisms and its relation to efficiency: Our experience from \$40 billion of combinatorial multi-attribute auctions, and recent theory

An extended abstract of an invited talk at the 2nd International Workshop on Computational Social Choice

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University

August 5, 2008

Abstract

A recent trend (especially in electronic commerce) is higher levels of expressiveness in the mechanisms that mediate interactions such as auctions, exchanges, catalog offers, voting systems, matching of peers, and so on. Participants can express their preferences in drastically greater detail than ever before. In many cases this trend is fueled by modern algorithms for winner determination that can handle the richer inputs.

But is more expressiveness always a good thing? What forms of expressiveness should be offered?

In this talk I will first report on our experience from over \$40 billion of combinatorial multi-attribute sourcing auctions. Then, I will present recent theory that ties the expressiveness of a mechanism to an upper bound on efficiency in a domain-independent way in private-information settings. Time permitting, I will also discuss theory and experiments on applying expressiveness to ad auctions, such as sponsored search and real-time banner ad auctions with temporal span and complex preferences.

1 Introduction

By carefully crafting mechanisms it is possible to design better auctions, exchanges, catalog offers, voting systems, and so on. A recent trend in the world—especially in electronic commerce—is a demand for higher levels of expressiveness in the mechanisms that mediate interactions such as the allocation of resources, matching of peers, or elicitation of privacy and security preferences.

The most famous expressive mechanism is a *combinatorial auction (CA)*, which allows participants to express valuations over *packages* of items. CAs have the recognized benefit of removing the exposure problems that bidders face when they have preferences over packages but in traditional auctions are allowed to submit bids on individual items only. CAs also have other acknowledged benefits.

Expressiveness also plays a key role in *multi-attribute* settings where the participants can express preferences over vectors of attributes of the item—or, more generally, of the outcome.

The trend toward expressiveness is also reflected in the richness of preference expression offered by businesses as diverse as matchmaking sites, sites like Amazon and Netflix, and services like Google’s AdSense. In Web 2.0 parlance, this demand for increasingly diverse offerings is called the Long Tail [1].

2 Our real-world experiences with expressive mechanisms in sourcing

In the first part of the talk, I will share some of my experiences from using expressiveness in practice. I started building winner determination algorithms for combinatorial auctions in 1997, and founded a company, CombineNet, Inc., in 2000 to field expressive mechanisms. Since then we have fielded over 500 expressive auctions. These auctions have been in the area of strategic sourcing, that is, the process by which large companies buy materials, products, services, and transportation from their suppliers, striking long-term contracts based on each auction.

Our auction designs, which we now call *expressive commerce*, hybridize and generalize both combinatorial and multi-attribute auctions [7, 9]. Expressive commerce 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 an 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 fastest custom tree search algorithms for solving it. We have hosted over \$40 billion of sourcing using the technology, and created over \$5 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.

We found that the traditional form of expressive bidding in CAs, package bidding (possibly with different forms of exclusivity constraints between bids), is a much too impoverished a bidding language to be usable in practice. In contrast, we found that there are a host of more compact and natural expressiveness constructs, and they are all used in concert in our auctions. These include various flexible forms of package bids, rich forms of conditional discount offers, various forms of discount schedules, side constraints, expressions of cost drivers, and multiattribute bidding [7].

In our events the bid taker can also express various forms of preferences and constraints. By conducting what-if analysis by changing these, the bid taker can form a

quantitative understanding of the tradeoffs available in the supply chain, such as cost versus multiple measures of practical implementability of the allocation, cost versus multiple measures of quality of the allocation, and cost versus multiple measures of long-term risk entailed by the allocation [7].

Furthermore, by allowing expressive offers over different combinations of the items to be sourced, the winner determination, as a side effect, ends up redesigning the supply chain. For example, in a sourcing 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, etc. [8]. 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.

An additional interesting aspect of bidding with cost drivers and alternates (e.g., using attributes) is that the winner determination algorithm not only decides who wins, but also ends up optimizing the configuration (setting of attributes) for each item, and the process by which each item is made.

3 Theory

Intuitively, one would think that increases in expressiveness would lead to more efficient mechanisms. That is also what the CombineNet experiences suggest. However, until now we have lacked a general theory that ties expressiveness and efficiency.

We developed a theory that ties the expressiveness of mechanisms to their efficiency in a domain-independent manner [3]. We introduce two new expressiveness measures, 1) *maximum impact dimension*, which captures the number of ways that an agent can impact the outcome, and 2) *shatterable outcome dimension*, which is based on the concept of *shattering* from computational learning theory. We derive an upper bound on the expected efficiency of any mechanism under its most efficient Nash equilibrium. Remarkably, it depends only on the mechanism’s expressiveness. We prove that the bound increases strictly as we allow more expressiveness. We also show that in some cases a small increase in expressiveness yields an arbitrarily large increase in the bound.

Finally, we study *channel-based* mechanisms. The restriction is that these mechanisms take expressions of value through channels from agents to outcomes, and select the outcome with the largest sum. (Channel-based mechanisms subsume most combinatorial and multi-attribute auctions, the Vickrey-Clarke-Groves mechanism, etc.) In this class, a natural measure of expressiveness is the number of channels allowed (this generalizes the k -wise dependence measure of expressiveness used in the combinatorial auction literature). We show that our domain-independent measures of expressiveness appropriately relate to the natural measure of expressiveness of channel-based mechanisms: the number of channels allowed. Using this bridge, our general results yield interesting implications. For example, any (channel-based) multi-item auction that does not allow rich combinatorial bids can be arbitrarily inefficient—unless agents

have no private information.

4 Applications to ad auctions and exchanges

Advertisement auctions and exchanges are relatively new forms of buying and selling ad space. They are an opportune next area of application for expressive mechanisms.

4.1 The case of an isolated sponsored search auction

Sponsored search auctions (the dispatch of typically textual ads in response to keyword-based web searches) account for tens of billions of dollars in revenue annually (e.g., to Google, Yahoo, and Microsoft) and are some of the fastest growing mechanisms on the Internet. However, the most frequent variant of these mechanisms does not allow bidders to offer a separate bid for each ad position, and is thus inexpressive on a fundamental level. Here we attempt to characterize the cost of this inexpressiveness [2]. We adapt the theoretical framework discussed in the previous section to show that the commonly used *generalized second price (GSP)* mechanism is arbitrarily inefficient for some distributions over agent preferences. We then describe a search technique that computes an upper bound on the expected efficiency of the GSP mechanism for any given distribution over agent preferences. We report the results of running our search technique on synthetic preference distributions. Our results demonstrate that the cost of inexpressiveness is most severe when agents have diverse preferences (such as having both brand advertisers and value advertisers in the auction) and relatively low profit margins. Our results also show that designating one or more positions as “premium” and soliciting an extra bid for these positions eliminates almost all of the inefficiency.

4.2 Highly expressive real-time ad auctions that span time

The prevalence and variety of online advertising in recent years has led to the development of an array of services for both advertisers and purveyors of media. Because matching an advertiser’s needs (demand) with a content provider’s properties (e.g., locations on displayed web pages) is a complex enterprise, often automated matching is used to match ad channels¹ with advertisers. One famous example is sponsored search. Internet auctions of traditional advertising (TV, radio, print) are also emerging (e.g., via companies like *Google* and *Spot Runner*). Auctions and exchanges for banner ads have also been established—e.g., *Right Media* (now part of *Yahoo!*) and *DoubleClick* (now part of *Google*)—although many banner ad bulk contracts are still manually negotiated.

There has been considerable research on developing auction mechanisms for allocating ad channels, with a focus on issues like auction design, charging schemes (e.g., per impression or per *click-through (CT)*), bidder strategies, and so on. However, attention has focused almost exclusively on improving single-period expressiveness, still with per-impression or per-CT prices. As has been well-documented in other auction

¹Here we use the word “channel” totally differently than in the “channel-based” mechanisms discussed earlier in this abstract. Here, each “channel” is a subset of supply such that no bid distinguishes between different forms of supply within the channel.

domains like sourcing, requiring bidders and bid takers to shoehorn their preferences into the impoverished language of per-item bids is usually undesirably restrictive.

Here we explore the use of *expressive bidding* for online banner ad auctions. For ease of presentation, we discuss banner ads, but the general principles and specific techniques we propose can be applied to other forms of online advertising (electronic auctions of TV and radio ads, sponsored search, etc.) as well.

In many domains, the value of a *set* of ads may not be an additive function of value of its individual elements. For instance, in an advertising campaign, *campaign-level* expressiveness is important. Advertisers may value particular *sequences* of ads, rather than individual ads per se. Efficiency (and revenue) maximization in such an environment demand that we allow bidders to express bids (propose *contracts*) on complex allocations, and that bid takers optimize over *sequences* of allocations to best match bidder preferences, in a way that cannot be accommodated using per-item bidding.

The key technical challenge for expressive ad auctions is optimization: determining the optimal allocation of ad channels to very large numbers of complex bids in real-time. This is further complicated by the stochastic nature of the domain—both *supply* (number of impressions or CTs) and *demand* (future bids) are uncertain—which suggests the need for online allocation.

To address these issues, we introduced the idea of an *optimize-and-dispatch* architecture [6] where an optimizer is run only every so often and it parameterizes a dispatcher that operates in real time. The optimizer can be run at fixed intervals, or based on any other trigger conditions, such as supply or demand significantly deviating from their projections. The framework can, in principle, handle any forms of expressive preferences as inputs, and we discuss several forms of expressiveness that are important in ad auctions, but which prior ad auction mechanisms inherently cannot support.

We recently implemented these ideas [4]. We model the problem as a Markov decision process (MDP), whose solution is approximated in several ways. First we perform optimization only periodically. Following the general optimize-and-dispatch framework, our optimization generates an on-line *dispatch policy* that assigns ad channels to advertisers in real-time. Our dispatch policies use the fractional assignment of (dynamically defined) channels to specific contracts. To approximate the optimization itself, we consider two approaches. The first is deterministic optimization using expectations of all random variables and exploiting our combinatorial optimization technology for winner determination in expressive markets [7]. We propose a second, sample-based approach derived from van Hentenryck and Bent’s [5] online model for stochastic optimization—but with novel adaptations to a continuous decision space. This approach leverages the deterministic winner determination technology, applying it to multiple possible future scenarios in order to form a dispatch policy. In both cases, periodic reoptimization is used to overcome the approximate nature of the methods. Our experiments demonstrate the benefits of expressive bidding for ad auctions over various per-item strategies, and the value of our stochastic optimization techniques.

References

- [1] Chris Anderson. *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion, July 2006.
- [2] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. The cost of inexpressiveness in advertisement auctions. In *Proceedings of the Fourth Workshop on Ad Auctions*, 2008.
- [3] Michael Benisch, Norman Sadeh, and Tuomas Sandholm. A theory of expressiveness in mechanisms. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 2008.
- [4] Craig Boutilier, David Parkes, Tuomas Sandholm, and William Walsh. Expressive banner ad auctions and model-based online optimization for clearing. In *Proceedings of National Conference on Artificial Intelligence (AAAI)*, 2008.
- [5] Pascal Van Hentenryck and Russell Bent. *Online Stochastic Combinatorial Optimization*. MIT Press, 2006.
- [6] David Parkes and Tuomas Sandholm. Optimize-and-dispatch architecture for expressive ad auctions. In *First Workshop on Sponsored Search Auctions, at the ACM Conference on Electronic Commerce*, Vancouver, BC, Canada, June 2005.
- [7] Tuomas Sandholm. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, 28(3):45–58, 2007.
- [8] Tuomas Sandholm, David Levine, Michael Concordia, Paul Martyn, Rick Hughes, Jim Jacobs, and Dennis Begg. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. *Interfaces*, 36(1):55–68, 2006.
- [9] Tuomas Sandholm and Subhash Suri. Side constraints and non-price attributes in markets. *Games and Economic Behavior*, 55:321–330, 2006. Extended version in IJCAI-2001 Workshop on Distributed Constraint Reasoning.