

# Computing in Mechanism Design\*

**Tuomas Sandholm**  
Computer Science Department  
Carnegie Mellon University  
Pittsburgh, PA 15213

## 1. Introduction

Computational issues in mechanism design are important, but have received insufficient research interest until recently. Limited computing hinders mechanism design in several ways, and presents deep strategic interactions between computing and incentives. On the bright side, the vast increase in computing power has enabled better mechanisms. Perhaps most interestingly, limited computing of the agents can be used as a tool to implement mechanisms that would not be implementable among computationally unlimited agents. This chapter briefly reviews some of the key ideas, with the goal of alerting the reader to the importance of these issues and hopefully spurring future research.

I will discuss computing by the *center*, such as an auction server or vote aggregator, in Section 2. Then, in Section 3, I will address the *agents'* computing, be they human or software.

## 2. Computing by the center

Computing by the center plays significant roles in mechanism design. In the following three subsections I will review three prominent directions.

### 2.1. Executing expressive mechanisms

As algorithms have advanced drastically and computing power has increased, it has become feasible to field mechanisms that were previously impractical. The most famous example is a *combinatorial auction (CA)*. In a CA, there are multiple distinguishable items for sale, and the bidders can submit bids on self-selected packages of the items. (Sometimes each bidder is also allowed to submit exclusivity constraints of different forms among his bids.) This increase in the expressiveness of the bids drastically reduces the strategic complexity that bidders face. For one, it removes 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 shift the computational burden from the bidders to the center. There is an associated gain because the center has all the information in hand to optimize while in traditional auctions the bidders only have estimated projected (probabilistic) information about how others will bid. Thus CAs yield more efficient allocations.

On the downside, the center's task of determining the winners in a CA (deciding which bids to accept so as to maximize the sum of the accepted bids' prices subject to not selling any item to more than one bid) is a complex combinatorial optimization problem, even without exclusivity constraints among bids. Three main approaches have been studied for solving it:

- *Optimal winner determination using some form of tree search.* For a review, see Sandholm (2006). The advantage is that the bidding language is not restricted and the optimal solution is found. The downside is that no optimal winner determination algorithm can run in polynomial time in the size of the problem instance in the worst case, because the problem is  $\mathcal{NP}$ -complete (Rothkopf et al., 1998). ( $\mathcal{NP}$ -complete problems are problems for which the fastest known algorithms take exponential time in the size of the problem instance in the worst case.  $\mathcal{P}$  is the class of easy problems solvable in polynomial time. The statement of winner determination not being solvable in polynomial time in the worst case relies on the usual assumption  $\mathcal{P} \neq \mathcal{NP}$ . This is an open question in complexity theory, but is widely believed to be true. If false, that would have sweeping implications throughout computer science.)

---

\*This work was funded by the National Science Foundation under ITR grant IIS-0427858, and a Sloan Foundation Fellowship. I thank Felix Brandt, Christina Fong, Joe Halpern, and David Parkes for helpful comments on drafts of this paper.

- *Approximate winner determination.* The advantage is that many approximation algorithms run in polynomial time in the size of the instance even in the worst case. For reviews of such algorithms, see Sandholm (2002a) and Lehmann et al. (2006). (Other suboptimal algorithms do not have such time guarantees, such as local search, stochastic local search, simulated annealing, genetic algorithms, and tabu search.) The downside is that the solution is sometimes far from optimal: no such algorithm can always find a solution that is within a factor

$$\min \left\{ \#bids^{1-\epsilon}, \#items^{\frac{1}{2}-\epsilon} \right\} \quad (1)$$

of optimal (Sandholm, 2002a). (This assumes  $\mathcal{ZPP} \neq \mathcal{NP}$ . It is widely believed that these two complexity classes are indeed unequal.) For example, with just 9 items for sale, no such algorithm can extract even 33% of the available revenue from the bids in the worst case. With 81 items, that drops to 11%.

- *Restricting the bidding language* so much that optimal (within the restricted language) winner determination can be conducted. For a review, see Müller (2006). For example, if each package bid is only allowed to include at most two items, then winners can be determined in worst-case polynomial time (Rothkopf et al., 1998). The downside is that bidders have to shoe-horn their preferences into a restricted bidding language; this gives rise to similar problems as in non-combinatorial mechanisms for multi-item auctions: exposure problems, need to speculate how others will bid, inefficient allocation, etc.

Truthful bidding can be made a dominant strategy by applying the *Vickrey-Clarke-Groves (VCG) mechanism* to a CA. Such incentive compatibility removes strategic complexity of the bidders. The mechanism works as follows. The optimal allocation is used, but the bidders do not pay their winning bids. Instead each bidder pays the amount of value he takes away from the others by taking some of the items. This value is measured as the difference between the others' winning bids' prices and what the others' winning bids' prices would have been had the agent not submitted any bids. This mechanism can be executed by determining the winners once overall, and once for each agent removed in turn. (This may be accomplishable with less computing. For example, in certain network auctions it can be done in the same asymptotic complexity as one winner determination (Hershberger and Suri, 2001).)

Very few canonical CAs have found their way to practice. However, auctions with richer bid expressiveness forms (that are more natural in the given application and more concise) and that support expressiveness also by the bid taker, have made a major breakthrough into practice (Sandholm, 2007; Bichler et al., 2006). This is sometimes called *expressive commerce* to distinguish it from vanilla CAs. The widest area of application is currently industrial sourcing. Tens of billions of dollars worth of materials, transportation, and services are being sourced annually using such mechanisms, yielding billions of dollars in efficiency improvements. The bidders' expressiveness forms include different forms of flexible package bids, conditional discounts, discount schedules, side constraints (such as capacity constraints), and often hundreds of cost drivers (e.g., fixed costs, variable costs, transshipment costs, and costs associated with changes). The item specifications can also be left partially open, and the bidders can specify some of the item attributes (delivery date, insurance terms, etc.) in alternate ways. The bid taker also specifies preferences and constraints. Winner determination then not only decides who wins what, but also automatically configures the items. In some of these events it also configures the supply chain several levels deep as a side effect. On the high end, such an auction can have tens of thousands of items (multiple units of each), millions of bids, and hundreds of thousands of side constraints. Expressive mechanisms have also been designed for settings beyond auctions, such as combinatorial exchanges, charity donations, and settings with externalities.

Basically all of the fielded expressive auctions use the simple pay-your-winning-bids pricing rule. There are numerous important reasons why few, if any, use the VCG mechanism. It can lead to low revenue. It is vulnerable to collusion. Bidders would not tell the truth because they do not want to reveal their cost structures which the auctioneer could exploit the next time when the auction is conducted, and so on (Sandholm, 2000; Rothkopf, 2007).

Basically all of the fielded expressive auctions use tree search for winner determination. In practice, modern tree search algorithms for the problem scale to the large and winners can be determined optimally. If winner determination were not done optimally in a CA, the VCG mechanism can lose its truth-dominance property (Sandholm, 2002b). In fact, any truthful suboptimal VCG-based mechanism for CAs is unreasonable in the sense that it sometimes does not allocate an item to a bidder even if he is the only bidder whose bids assign non-zero value to that item (Nisan and Ronen, 2000).

## 2.2. Algorithmic mechanism design

Motivated by the worry that some instances of  $\mathcal{NP}$ -hard problems may not be solvable within reasonable time, a common research direction in theory of computing is approximation algorithms. They trade off solution quality for a guarantee that even in the worst case, the algorithm runs in polynomial time in the size of the input.

Analogously, Nisan and Ronen (2001) proposed *algorithmic mechanism design*: designing approximately optimal mechanisms that take the center a polynomial number of computing steps even in the worst case. However, this is more difficult than designing approximately optimal algorithms because the mechanism has to motivate the agents to tell the truth.

Lehmann et al. (2002) studied this for CAs with single-minded bidders (each bidder being only interested in one specific package of items). They present a fast greedy algorithm that guarantees a solution within a factor  $\sqrt{\#items}$  of optimal. They show that the algorithm is not incentive compatible with VCG pricing, but is with their custom pricing scheme. They also identify sufficient conditions for any (approximate) mechanism to be incentive compatible. See also Kfir-Dahav et al. (2000). There has been substantial follow-on work on subclasses of single-minded CAs.

Lavi and Swamy (2005) developed a technique for a range of packing problems with which any  $k$ -approximation algorithm (i.e., algorithm that guarantees that the solution is within a factor  $k$  of optimal) that also bounds the integrality gap of the linear programming (LP) relaxation of the problem by  $k$  can be used to construct a  $k$ -approximation mechanism. The LP solution, scaled down by  $k$ , can be represented as a convex combination of integer solutions, and viewing this convex combination as specifying a probability distribution over integer solutions begets a VCG-based randomized mechanism that is truthful in expectation. For CAs with general valuations, this yields an  $O(\sqrt{\#items})$ -approximate mechanism.

In a different direction, several mechanisms have been proposed where the agents can help the center find better outcomes. This is done either by giving the agents the information to do the center's computing (Banks et al., 1989; Land et al., 2006; Parkes and Shneidman, 2004), or by allowing the agents to change what they told the mechanism based on the mechanism's output and potentially also based on what other agents told the mechanism (Nisan and Ronen, 2000). In VCG-based mechanisms, an agent benefits from lying only if the lie causes the mechanism to find an outcome that is better overall.

## 2.3. Automated mechanism design

Conitzer and Sandholm (2002) proposed the idea of *automated mechanism design*: having a computer, rather than a human, design the mechanism. Because human effort is eliminated, this enables custom design of mechanisms for every setting. (The setting can be described by the agents' (discretized) type spaces, the designer's prior over types, the desired notion of incentive compatibility (e.g., dominant strategies vs. Bayes-Nash implementation), the desired notion of participation constraints (e.g., *ex interim*, *ex post*, or none), whether payments are allowed, and whether the mechanism is allowed to use randomization.) This can yield better mechanisms for previously studied settings because the mechanism is designed for the specific setting rather than a class of settings. It can also be used for settings not previously studied in mechanism design.

For almost all natural (linear) objectives, all variants of the design problem are  $\mathcal{NP}$ -complete if the mechanism is not allowed to use randomization, but randomized mechanisms can be constructed for all these settings in polynomial time using linear programming. Custom algorithms have been developed for some problems in each of these two categories. (Even the latter category warrants research. While the linear program is polynomial in the size of the input, the input itself can be exponential in the number of agents.) Structured representations of the problem can also make the design process drastically faster.

Beyond the general setting, automated mechanism design has been applied to specific settings, such as creating revenue-maximizing CAs (without the need to discretize types) (Likhodedov and Sandholm, 2005) (a recognized problem that eludes analytical characterization; even the 2-item case is open), reputation systems (Jurca and Faltings, 2006), safe exchange mechanisms (Sandholm and Ferrandon, 2000), and supply chain settings (Vorobeychik et al., 2006). Automated mechanism design software has recently also been adopted by several mechanism design theoreticians to speed up their research.

It turns out that even *multistage mechanisms* can be designed automatically (Sandholm et al., 2007). Furthermore, automated mechanism design has been applied to the design of *online mechanisms* (Hajiaghayi et al., 2007), that is, mechanisms that execute while the world changes—e.g., agents enter and exit the system.

### 3. Computing by the agents

I will now move to discussing computing by the agents.

#### 3.1. Mechanisms that are hard to manipulate

This section demonstrates that one can use the fact that agents are computationally limited to achieve things that are not achievable via any mechanism among perfectly rational agents.

A seminal negative result, the *Gibbard-Satterthwaite theorem*, states that if there are three or more candidates, then in any nondictatorial voting scheme, there are candidate rankings of the other voters, and preferences of the agent, under which the agent is better off voting manipulatively than truthfully. One avenue around this impossibility is to construct desirable general nondictatorial voting protocols under which *finding* a beneficial manipulation is prohibitively hard computationally.

There are two natural alternative goals of manipulation. In *constructive manipulation*, the manipulator tries to find an order of candidates that he can reveal so that his favorite candidate wins. In *destructive manipulation*, the manipulator tries to find an order of candidates that he can reveal so that his hated candidate does not win. These are special cases of the utility-theoretic notion of improving one’s utility, so the hardness results carry over to the usual utility-theoretic setting.

Unfortunately, finding a constructive manipulation is easy (in  $\mathcal{P}$ ) for the *plurality*, *Borda*, and *maximin* voting rules (Bartholdi et al., 1989), which are commonly used. On the bright side, constructive manipulation of the *Single Transferable Vote (STV)* protocol is  $\mathcal{NP}$ -hard (Bartholdi and Orlin, 1991) (as is manipulation of the is the *Second Order Copeland* protocol (Bartholdi et al., 1989), but that hardness is driven solely by the tie-breaking rule). Even better, there is a systematic methodology for slightly tweaking the voting protocols that are easy to manipulate, so that they become hard to manipulate (Conitzer and Sandholm, 2003). Specifically, before the original protocol is executed, one pairwise elimination round is executed among the candidates, and only the winning candidates survive to the original protocol. This makes the protocols  $\mathcal{NP}$ -hard,  $\#\mathcal{P}$ -hard ( $\#\mathcal{P}$ -hard problems are at least as hard as counting the number of solutions to a problem in  $\mathcal{P}$ ), or even  $\mathcal{PSPACE}$ -hard ( $\mathcal{PSPACE}$ -hard problems are at least as hard as any problem that can be solved using a polynomial amount of memory) to manipulate constructively, depending on whether the schedule of the preround is determined before the votes are collected, randomly after the votes are collected, or the scheduling and the vote collecting are carefully interleaved, respectively.

All of the hardness results of the previous paragraph rely on both the number of voters and the number of candidates growing. The number of candidates can be large in some domains, for example when voting over task or resource allocations. However, in other elections—such as presidential elections—the number of candidates is small. If the number of candidates is a constant, both constructive and destructive manipulation are easy (in  $\mathcal{P}$ ), regardless of the number of voters (Conitzer et al., 2007). This holds even if the voters are weighted, or if a coalition of voters tries to manipulate. On the bright side, when a coalition of weighted voters tries to manipulate, complexity can arise even for a constant number of candidates, see Table 1 (Conitzer et al., 2007). Another lesson from that table is that randomizing over instantiations of the mechanism (such as schedules of a *cup*) can be used to make manipulation hard.

Number of candidates:	2	3	4,5,6	$\geq 7$
<i>Borda</i>	$\mathcal{P}$	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete
<i>veto</i>	$\mathcal{P}$	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete
<i>STV</i>	$\mathcal{P}$	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete
<i>plurality with runoff</i>	$\mathcal{P}$	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete
<i>Copeland</i>	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete
<i>maximin</i>	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{NP}$ -complete	$\mathcal{NP}$ -complete
<i>randomized cup</i>	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{NP}$ -complete
<i>cup</i>	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{P}$
<i>plurality</i>	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{P}$	$\mathcal{P}$

Number of candidates:	2	$\geq 3$
<i>STV</i>	$\mathcal{P}$	$\mathcal{NP}$ -complete
<i>plurality with runoff</i>	$\mathcal{P}$	$\mathcal{NP}$ -complete
<i>randomized cup</i>	$\mathcal{P}$	?
<i>Borda</i>	$\mathcal{P}$	$\mathcal{P}$
<i>veto</i>	$\mathcal{P}$	$\mathcal{P}$
<i>Copeland</i>	$\mathcal{P}$	$\mathcal{P}$
<i>maximin</i>	$\mathcal{P}$	$\mathcal{P}$
<i>cup</i>	$\mathcal{P}$	$\mathcal{P}$
<i>plurality</i>	$\mathcal{P}$	$\mathcal{P}$

Table 1: *Left: Complexity of constructive weighted coalitional manipulation. Right: Complexity of destructive weighted coalitional manipulation.*

As usual in computer science, the results are worst-case hardness. Unfortunately, under weak assumptions on the preference distribution and voting rule, most instances of any voting rule are easy to manipulate (Conitzer and Sandholm, 2006).

All of the hardness results discussed above hold even if the manipulators know the nonmanipulators’ votes exactly. Under weak assumptions, if weighted coalitional manipulation with complete information about the others’ votes is

hard in some voting protocol, then individual and unweighted manipulation is hard when there is uncertainty about the others' votes (Conitzer et al., 2007).

### 3.2. Non-truth-promoting mechanisms

A challenging issue is that even if it is prohibitively hard to find a beneficial manipulation, the agents might not tell the truth. For example, an agent might take a chance that he will do better with a lie. The following result shows that, nevertheless, mechanism design can be improved by making the agents face complexity. (This is one reason why computational issues can render the *revelation principle* inapplicable. One of the things the principle says is that for any non-truth-promoting mechanism it is possible to construct an incentive-compatible mechanism that is at least as good. The theorem below challenges this.)

**Theorem 1** (Conitzer and Sandholm, 2004) *Suppose the center is trying to maximize social welfare, and neither payments nor randomization is allowed. Then, even with just two agents (one of whom does not even report a type, so dominant strategy implementation and Bayes-Nash implementation coincide), there exists a family of preference aggregation settings such that:*

- *the execution of any optimal incentive-compatible mechanism is  $\mathcal{NP}$ -complete for the center, and*
- *there exists a non-incentive-compatible mechanism which 1) requires the center to carry out only polynomial computation, and 2) makes finding any beneficial insincere revelation  $\mathcal{NP}$ -complete for the type-reporting agent. Additionally, if the type-reporting agent manages to find a beneficial insincere revelation, or no beneficial insincere revelation exists, the social welfare of the outcome is identical to the social welfare that would be produced by any optimal incentive-compatible mechanism. Finally, if the type-reporting agent does not manage to find a beneficial insincere revelation where one exists, the **social welfare of the outcome is strictly greater than the social welfare that would be produced by any optimal incentive-compatible mechanism.***

An analogous theorem holds if, instead of counting computational steps, we count calls to a commonly accessible oracle which, when supplied with an agent, that agent's type, and an outcome, returns a utility value for that agent.

### 3.3. Preference (valuation) determination via computing or information acquisition

In many (auction) settings, even determining one's valuation for an item (or a bundle of items) is complex. For example when bidding for trucking lanes (tasks), this involves solving two  $\mathcal{NP}$ -complete local planning problems: the vehicle routing problem with the new lanes of the bundle and the problem without them (Sandholm, 1993). The difference in the costs of those two local plans is the cost (valuation) of taking on the new lanes.

In these types of settings, the *revelation principle* applies only in a trivial way: the agents report their data and optimization models to the center, and the center does the computation for them. It stands to reason that in many applications the center would not want to take on that burden, in which case such extreme direct mechanisms are not an option. Therefore, I will now focus on mechanisms where the agents report valuations to the center, as in traditional auctions.

Bidders usually have limited computing and time, so they cannot exactly evaluate all (or even any) bundles— at least not without cost. This leads to a host of interesting issues where computing and incentives are intimately intertwined.

For example, in a 1-object auction, should a bidder evaluate the object if there is a cost to doing so? It turns out that the Vickrey auction loses its dominant-strategy property: whether or not the bidder should pay the evaluation cost depends on the other bidders' valuations (Sandholm, 2000).

If a bidder has the opportunity to *approximate his valuation to different degrees*, how much computing time should the bidder spend on refining its valuation? If there are multiple items for sale, how much computing time should the bidder allocate on different bundles? A bidder may even allocate some computing time to evaluate other bidders' valuations so as to be able to bid more strategically; this is called *strategic computing*.

To answer these questions, Larson and Sandholm (2001) developed a deliberation control method called a *performance profile tree* for projecting how an anytime algorithm (i.e., an algorithm that has an answer available at any time, but where the quality of the answer improves the more computing time the algorithm is allocated) will change the valuation if additional computing is allocated toward refining (or improving) it. This deliberation control method applies to any anytime algorithm. Unlike earlier deliberation control methods for anytime algorithms, the performance

profile tree is a *fully normative model of bounded rationality*: it takes into account all the information that an agent can use to make its deliberation control decisions. This is necessary in the game-theoretic context; otherwise a strategic agent could take into account some information that the model does not.

Using this deliberation control method, the auction can be modeled as a game where the agents’ strategy spaces include computing actions. At every point, each agent can decide on which bundle to allocate its next step of computing as a function of the agent’s computing results so far (and in open-cry auction format also the others’ bids observed so far). At every point, the agent can also decide to submit bids. One can then solve this for equilibrium: each agent’s (deliberation and bidding) strategy is a best-response to the others’ strategies. This is called *deliberation equilibrium*.

This notion, and the performance profile tree, apply not only to computational actions but also to information gathering actions for determining valuations. (In contrast, most of the literature on information acquisition in auctions does not take into account that valuations can be determined to different degrees and that an agent may want to invest effort to determine others’ valuations as well—even in private-value settings.)

Table 2 shows in which settings strategic computing can and cannot occur in deliberation equilibrium. This depends on the auction mechanism. Interestingly, it also depends on whether the agent has limited computing (e.g., owning a desktop computer that the agent can use until the auction’s deadline) or costly computing (e.g., being able to buy any amount of supercomputer time where each cycle comes at a cost).

	Auction mechanism	Speculation by perfectly rational agents?	Strategic computing?	
			Limited computing	<b>Costly computing</b>
Single item	First price	yes	yes	yes
	Dutch	yes	yes	yes
	English	no	<b>no</b>	<b>yes</b>
	Vickrey	no	<b>no</b>	<b>yes</b>
Multiple items	First price	yes	yes	yes
	VCG	no	<b>yes</b>	<b>yes</b>

Table 2: *Can strategic computing occur in deliberation equilibrium? The most interesting results are in bold. As a benchmark from classical auction theory, the table also shows whether or not perfectly rational agents, that can determine their valuations instantly without cost, would benefit from considering each others’ valuations when deciding how to bid.*

The notion of deliberation equilibrium can also be used as the basis for designing new mechanism which hopefully would work well among agents whose computing is costly or limited. Unfortunately, there is an impossibility (Larson and Sandholm, 2005): there exists no mechanisms that is *sensitive* (the outcome is affected by each agent’s strategy), *preference formation independent* (does not do the computations for the agents; the agents report valuations), *non-misleading* (no agent acts in a way that causes others to believe his true type has zero probability), and *deliberation-proof* (no strategic computing occurs in equilibrium, i.e., agents compute only on their own problems). Current work involves designing mechanisms that take part in preference formation in a limited way: agents report their performance profile trees to the center, which then coordinates the deliberations incrementally as agents report deliberation results. Current research also includes designing mechanisms where strategic computing occurs but its wastefulness is limited.

### 3.3.1. Preference elicitation by the center

To reduce the agents’ preference determination effort, Conen and Sandholm (2001) proposed a framework where the center (aka *elicitor*) explicitly elicits preference information from the agents incrementally on an as-needed basis by posing queries to the agents. The center thereby builds a model of the agents’ preferences, and decides what to ask, and from which agent, based on this model. Usually the process can be terminated with the provably correct outcome while requiring only a small portion of the agents’ preferences to be determined. Multistage mechanisms can yield up to exponential savings in preference determination and communication effort the agents need to go through compared to single-stage mechanisms (Conitzer and Sandholm, 2004).

The explicit preference elicitation framework was originally proposed for CAs (but the approach has since been used for other settings as well, such as voting). For general valuations, an exponential number of bits in the number of items for sale has to be communicated in the worst case no matter what queries are used (Nisan and Segal, 2006). However, experimentally only a small fraction of the preference information needs to be elicited before the provably optimal solution is found. Furthermore, for valuations that have certain types of structure, even the worst-case number of queries needed is small. Research has also been done on the relative power of different query types.

If enough information is elicited to also determine the VCG payments, and these are the payments charged to the bidders, answering the elicitor's queries truthfully is an *ex post* equilibrium (a strengthening of Nash equilibrium that does not rely on priors). (This assumes there is no explicit cost or limit to valuation determination; mechanisms have also been designed for settings where there is an explicit cost (Larson, 2006).). This holds even if the agents are allowed to answer queries that the elicitor did not ask (e.g., queries that are easy for the agent to answer and which the agent thinks will significantly advance the elicitation process). We thus have a *pull-push mechanism* where both the center and the agents guide the preference elicitation (and preference determination/refinement by the agents). For a review, see Sandholm and Boutilier (2006). Ascending (combinatorial) auctions are an earlier special case, and have limited power compared to the general framework (Blumrosen and Nisan, 2005).

Preference elicitation can sometimes be computationally complex for the center. It can be complex to intelligently decide what to ask next, and from whom. It can also be complex to determine whether enough information has been elicited to determine the optimal outcome. Even if the elicitor knows that enough has been elicited, it can be complex to determine the outcome—e.g., allocation of items to bidders in some CAs.

### 3.4. Distributed (center-free) mechanisms

Computer scientists often have a preference for distributed applications that do not have any centralized coordination point (center). Depending on the application, the reasons for this preference may include avoiding a single vulnerable point of failure, distributing the computing effort (for computational efficiency or because the data is inherently distributed), and enhancing privacy. The preference carries over from traditional computer science applications to different forms of negotiation systems—for example, see Sandholm (1993) for an early distributed automated negotiation system for software agents.

Feigenbaum et al. (2005) have studied lowest-cost inter-domain routing on the Internet, modifying a distributed protocol so that the agents (routing domains) are motivated to report their true costs and the solution is found with minimal message passing. For a review of some other research topics in this space, see Feigenbaum and Shenker (2002).

One can go further by taking into account that agents might not choose to follow the prescribed protocol. They may cheat not only on information-revelation actions, but also on message-passing and computational actions. Despite computation actions not being observable by others, an agent can be motivated to compute as prescribed by tasking at least one other agent with the same computation, and comparing the results (Sandholm et al., 1999). Careful problem partitioning can also be used to achieve the same outcome without redundancy by only requiring agents to perform computing and message passing tasks that are in their own interest (Parkes and Shneidman, 2004). Shneidman and Parkes (2004) propose a general proof technique and instantiate it to provide a non-manipulable protocol for inter-domain routing. Monderer and Tennenholtz (1999) develop protocols for 1-item auctions executed among agents on a communication network. The protocols motivate the agents to correctly reveal preferences and communicate. For the setting where agents with private utility functions have to agree on variable assignments subject to side constraints (e.g., meeting scheduling), Petcu et al. (2006) developed a VCG-based distributed optimization protocol that finds the social welfare maximizing allocation and each agent is motivated to follow the protocol in terms of all three types of action. The only centralized party needed is a bank that can extract payments from the agents.

Cryptography is a powerful tool for achieving privacy when trying to execute a mechanism in a distributed way without a center, using private communication channels among the agents. Consider first the setting with passive adversaries, that is, agents that faithfully execute the specified distributed communication protocol, but who try to infer (at least something about) some agent's private information.

- If agents are computationally limited, e.g., they are assumed to be unable to factor large numbers, then arbitrary functions can be computed while guaranteeing that each agent maintains his privacy (except, of course, to the extent that the answer of the computation says something about the inputs) (Goldreich et al., 1987). Thus the desire for privacy does not constrain what social choice functions can be implemented.
- In contrast, only very limited social choice functions can be computed privately among computationally unlimited agents. For example, when there are just two alternatives, every monotonic, non-dictatorial social choice function that can be privately computed is constant (Brandt and Sandholm, 2005). With special structure in the preferences, this impossibility can sometimes be avoided. For example, with the standard model of quasilinear utility, first-price auctions can be implemented privately; second-price (Vickrey) auctions with more than two bidders cannot (Brandt and Sandholm, 2004).

A more general model is that of active adversaries who can execute the distributed communication protocol unfaithfully in a coordinated way. A more game-theoretic model is that of rational adversaries that are not passive, but not malicious either. For a brief overview of such work, see Section 4 of the Palgrave chapter “Computer Science and Game Theory: A Brief Survey”.

## References

- Banks, Jeffrey S, John Ledyard, David Porter. 1989. Allocating uncertain and unresponsive resources: An experimental approach. *RAND Journal of Economics* **20** 1–25.
- Bartholdi, John, III, James Orlin. 1991. Single transferable vote resists strategic voting. *Social Choice and Welfare* **8** 341–354.
- Bartholdi, John, III, Craig Tovey, Michael Trick. 1989. The computational difficulty of manipulating an election. *Social Choice and Welfare* **6** 227–241.
- Bichler, Martin, Andrew Davenport, Gail Hohner, Jayant Kalagnanam. 2006. Industrial procurement auctions. Peter Cramton, Yoav Shoham, Richard Steinberg, eds., *Combinatorial Auctions*. MIT Press, 593–612. Chapter 23.
- Blumrosen, Liad, Noam Nisan. 2005. On the computational power of iterative auctions. *Proc. ACM Conference on Electronic Commerce*.
- Brandt, Felix, Tuomas Sandholm. 2004. (Im)possibility of unconditionally privacy-preserving auctions. *Proc. International Conference on Autonomous Agents and Multi-Agent Systems*.
- Brandt, Felix, Tuomas Sandholm. 2005. Unconditional privacy in social choice. *Proc. Conference on Theoretical Aspects of Rationality and Knowledge*.
- Conen, Wolfram, Tuomas Sandholm. 2001. Preference elicitation in combinatorial auctions: Extended abstract. *Proc. ACM Conference on Electronic Commerce*. More detailed description of algorithmic aspects in Proc. IJCAI-01 Workshop on Economic Agents, Models, and Mechanisms.
- Conitzer, Vincent, Tuomas Sandholm. 2002. Complexity of mechanism design. *Proc. Conference on Uncertainty in Artificial Intelligence*.
- Conitzer, Vincent, Tuomas Sandholm. 2003. Universal voting protocol tweaks to make manipulation hard. *Proc. International Joint Conference on Artificial Intelligence*.
- Conitzer, Vincent, Tuomas Sandholm. 2004. Computational criticisms of the revelation principle. *Conference on Logic and the Foundations of Game and Decision Theory*. Earlier versions: AMEC-03, EC-04.
- Conitzer, Vincent, Tuomas Sandholm. 2006. Nonexistence of voting rules that are usually hard to manipulate. *Proc. National Conference on Artificial Intelligence*.
- Conitzer, Vincent, Tuomas Sandholm, Jérôme Lang. 2007. When are elections with few candidates hard to manipulate? *Journal of the ACM* To appear.
- Feigenbaum, Joan, Christos Papadimitriou, Rahul Sami, Scott Shenker. 2005. A BGP-based mechanism for lowest-cost routing. *Distributed Computing* **18** 61–72.
- Feigenbaum, Joan, Scott Shenker. 2002. Distributed algorithmic mechanism design: Recent results and future directions. *Proc. International Workshop on Discrete Algorithms and Methods for Mobile Computing and Communications*.
- Goldreich, O., S. Micali, A. Wigderson. 1987. How to play any mental game or a completeness theorem for protocols with honest majority. *Proc. Symposium on Theory of Computing*.
- Hajiaghayi, MohammadTaghi, Robert Kleinberg, Tuomas Sandholm. 2007. Automated online mechanism design and prophet inequalities. *Proc. National Conference on Artificial Intelligence*.



- Hershberger, John, Subhash Suri. 2001. Vickrey prices and shortest paths: What is an edge worth? *Proc. Symposium on Foundations of Computer Science*.
- Jurca, Radu, Boi Faltings. 2006. Minimum payments that reward honest reputation feedback. *Proc. ACM Conference on Electronic Commerce*.
- Kfir-Dahav, Noa, Dov Monderer, Moshe Tennenholtz. 2000. Mechanism design for resource bounded agents. *Proc. International Conference on Multi-Agent Systems*.
- Land, Ailsa, Susan Powell, Richard Steinberg. 2006. PAUSE: A computationally tractable combinatorial auction. Peter Cramton, Yoav Shoham, Richard Steinberg, eds., *Combinatorial Auctions*, chap. 6. MIT Press.
- Larson, Kate. 2006. Reducing costly information acquisition in auctions. *Proc. Autonomous Agents and Multi-Agent Systems*.
- Larson, Kate, Tuomas Sandholm. 2001. Costly valuation computation in auctions. *Proc. Theoretical Aspects of Rationality and Knowledge*.
- Larson, Kate, Tuomas Sandholm. 2005. Mechanism design and deliberative agents. *Proc. International Conference on Autonomous Agents and Multi-Agent Systems*.
- Lavi, Ron, Chaitanya Swamy. 2005. Truthful and near-optimal mechanism design via linear programming. *Proc. Symposium on Foundations of Computer Science*.
- Lehmann, Daniel, Rudolf Müller, Tuomas Sandholm. 2006. The winner determination problem. Peter Cramton, Yoav Shoham, Richard Steinberg, eds., *Combinatorial Auctions*. MIT Press, 297–317. Chapter 12.
- Lehmann, Daniel, Lidian Ita O’Callaghan, Yoav Shoham. 2002. Truth revelation in rapid, approximately efficient combinatorial auctions. *Journal of the ACM* **49** 577–602.
- Likhodedov, Anton, Tuomas Sandholm. 2005. Approximating revenue-maximizing combinatorial auctions. *Proc. National Conference on Artificial Intelligence*.
- Monderer, Dov, Moshe Tennenholtz. 1999. Distributed games: From mechanisms to protocols. *Proc. National Conference on Artificial Intelligence*.
- Müller, Rudolf. 2006. Tractable cases of the winner determination problem. Peter Cramton, Yoav Shoham, Richard Steinberg, eds., *Combinatorial Auctions*. MIT Press, 319–336. Chapter 13.
- Nisan, Noam, Amir Ronen. 2000. Computationally feasible VCG mechanisms. *Proc. ACM Conference on Electronic Commerce*.
- Nisan, Noam, Amir Ronen. 2001. Algorithmic mechanism design. *Games and Economic Behavior* **35** 166–196.
- Nisan, Noam, Ilya Segal. 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* **129** 192–224.
- Parkes, David, Jeffrey Shneidman. 2004. Distributed implementations of generalized Vickrey-Clarke-Groves auctions. *Proc. International Conference on Autonomous Agents and Multi-Agent Systems*.
- Petcu, Adrian, Boi Faltings, David Parkes. 2006. MDPOP: Faithful distributed implementation of efficient social choice problems. *Proc. International Conference on Autonomous Agents and Multi-Agent Systems*.
- Rothkopf, Michael. 2007. Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. *Operations Research* **55** 191–197.
- Rothkopf, Michael, Aleksandar Pekeč, Ronald Harstad. 1998. Computationally manageable combinatorial auctions. *Management Science* **44** 1131–1147.
- Sandholm, Tuomas. 1993. An implementation of the contract net protocol based on marginal cost calculations. *Proc. National Conference on Artificial Intelligence*.

- Sandholm, Tuomas. 2000. Issues in computational Vickrey auctions. *International Journal of Electronic Commerce* **4** 107–129. Early version in ICMAS-96.
- Sandholm, Tuomas. 2002a. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence* **135** 1–54. Earlier versions: ICE-98 keynote, Washington U. tech report WUCS-99-01 Jan. 1999, IJCAI-99.
- Sandholm, Tuomas. 2002b. eMediator: A next generation electronic commerce server. *Computational Intelligence* **18** 656–676. Earlier versions: Washington U. tech report WU-CS-99-02 Jan. 1999, AAAI-99 Workshop on AI in Ecommerce, AGENTS-00.
- Sandholm, Tuomas. 2006. Optimal winner determination algorithms. Peter Cramton, Yoav Shoham, Richard Steinberg, eds., *Combinatorial Auctions*. MIT Press, 337–368. Chapter 14.
- Sandholm, Tuomas. 2007. Expressive commerce and its application to sourcing: How we conducted \$35 billion of generalized combinatorial auctions. *AI Magazine*, To appear.
- Sandholm, Tuomas, Craig Boutilier. 2006. Preference elicitation in combinatorial auctions. Peter Cramton, Yoav Shoham, Richard Steinberg, eds., *Combinatorial Auctions*. MIT Press, 233–263. Chapter 10.
- Sandholm, Tuomas, Vincent Conitzer, Craig Boutilier. 2007. Automated design of multistage mechanisms. *Proc. International Joint Conference on Artificial Intelligence*.
- Sandholm, Tuomas, Vincent Ferrandon. 2000. Safe exchange planner. *Proc. International Conference on Multi-Agent Systems*.
- Sandholm, Tuomas, Kate Larson, Martin Andersson, Onn Shehory, Fernando Tohmé. 1999. Coalition structure generation with worst case guarantees. *Artificial Intelligence* **111** 209–238.
- Shneidman, Jeffrey, David C. Parkes. 2004. Specification faithfulness in networks with rational nodes. *Proc. ACM Symposium on Principles of Distributed Computing*.
- Vorobeychik, Yevgeniy, Christopher Kiekintveld, Michael Wellman. 2006. Empirical mechanism design: Methods, with application to a supply chain scenario. *Proc. ACM Conference on Electronic Commerce*.