Beyond VCG: Frugality of Truthful Mechanisms

Anna R. Karlin[*]
University of Washington, Seattle
karlin@cs.washington.edu

David Kempe[†]
University of Southern California
dkempe@usc.edu

Tami Tamir[‡]
The Interdisciplinary Center, Herzliya, Israel
tami@idc.ac.il

Abstract

We study truthful mechanisms for auctions in which the auctioneer is trying to hire a team of agents to perform a complex task, and paying them for their work. As common in the field of mechanism design, we assume that the agents are selfish and will act in such a way as to maximize their profit, which in particular may include misrepresenting their true incurred cost. Our first contribution is a new and natural definition of the frugality ratio of a mechanism, measuring the amount by which a mechanism "overpays", and extending previous definitions to all monopoly-free set systems.

After reexamining several known results in light of this new definition, we proceed to study in detail shortest path auctions and "r-out-of-k sets" auctions. We show that when individual set systems (e.g., graphs) are considered instead of worst cases over all instances, these problems exhibit a rich structure, and the performance of mechanisms may be vastly different. In particular, we show that the well-known VCG mechanism may be far from optimal in these settings, and we propose and analyze a mechanism that is always within a constant factor of optimal.

1 Introduction

The design of protocols for resource allocation and electronic commerce among parties with diverse and selfish interests has spawned a great deal of recent research at the boundary between economics, game theory, and theoretical computer science. In many settings, a natural way to assign resources to, or obtain goods or services from, such selfish parties is by means of *auctions*, in which the parties submit bids to the auctioneer, who then chooses one or more winners and purchases their services (or sells them resources).

The case of buying or selling a single item at a time has a long history of study in economics [15, 14], and is well understood. The problem becomes more intriguing, however, when the auctioneer is trying to buy or sell a set of items, interacting with multiple parties. Several recent papers have studied the resulting auctions and their properties. We adopt here the general framework of *hiring a team of agents* [4, 27, 10]: The auctioneer is intent on hiring a team of agents to perform a complex task. Each agent i can perform a simple task at some cost c_i known only to himself. Based on the agents' bids b_i, the auctioneer must select a *feasible set* — a set of agents whose combined skills are sufficient to perform the complex task — and pay each selected agent individually some amount p_i. In the absence of the agents' costs and bids, the problem is therefore defined entirely by the *set system* of feasible sets.

Two special cases of this general setting have been studied extensively in the past: (i) In a *path auction* [22, 4, 8], the agents own edges of a known graph, and the auctioneer wants to purchase a path between two given nodes s, t. (ii) In a *minimum spanning tree (MST) auction* [27, 5, 10], the agents again own edges of a graph, and the auctioneer wants to purchase a spanning tree.

As mentioned above, agents in the real world will be selfish, and thus take whichever action benefits them most. In particular, they may report a value different from their true cost c_i to the auctioneer, in order to increase the payments they receive. The area of *mechanism design* [17, 22, 25] studies the design of auctions so that no agent, motivated only by self-interest, has an incentive to cheat.

A desirable property for mechanisms is that it be in each agent's best interest to report his actual cost c_i as his bid, no matter how other agents bid. This property of *truthfulness* obviates the need for agents to perform complex computa-

[*] Work supported by NSF CCR-0105406.

[†] Work done while the author was at the University of Washington, supported by an NSF Postdoctoral Fellowship.

[‡] Work done while the author was at the University of Washington.

tions or gather data about their competition, and at the same time simplifies the design and analysis of auction protocols, as there is no need for assumptions about agents' knowledge of each other or the distributions of their costs. In addition, the revelation principle ([17], p. 871) shows that there is no loss of generality in restricting attention to truthful mechanisms, so long as mechanisms are required to have *dominant strategies*; i.e., each agent has an optimal bid $\hat{b}_i$ based only on his own cost c_i, and independent of all other agents' bids.

For single-item auctions, the classical truthful mechanism is the Vickrey or Second-Price Auction [28], in which the lowest bidder is awarded the contract and is paid the second-lowest bid;[1] its generalization to the problem of hiring a team of agents and many other problems is called the VCG mechanism [28, 6, 12]. In the VCG mechanism, the feasible set S selected is always the one with lowest total cost $b(S)$ according to the agents' bids; each agent is then paid the highest amount it could have bid to still be part of the winning feasible set, all other agents' bids remaining the same. The VCG mechanism is truthful, and hence the winning set is in fact a cheapest feasible solution with respect to the true cost. However, the payments made by the VCG mechanism can be significantly greater than the cost of achieving truthfulness.

1.1 Frugality

In one traditional economical view, this overpayment has not been a major concern, as the auctioneer was assumed to be a central authority with a common social good in mind. Hence, any payments that were not directly applied toward actual costs increased the "social welfare", independently of whether they were in the hands of the auctioneer or the agents. An alternative goal, however, is for the auctioneer to maximize his own benefit, which includes hiring the team of agents as cheaply as possible. Hence, analyzing the *frugality* [4, 27, 10] of a mechanism — the amount by which it overpays — becomes an important aspect of mechanism design. The fundamental issues are:

- How do we design truthful mechanisms that are as frugal as possible?

- How costly is the restriction to truthfulness?

Here, we build on research initiated by Archer and Tardos, Talwar, and others [3, 4, 27, 10, 22, 8, 5, 11], and study the frugality of mechanisms for several classes of the "hiring a team" problem.

When analyzing the overpayment of a mechanism, the first important question is "Overpayment compared to what?". The first naïve approach would be to compare to the cost of an actual cheapest solution: however, even for the simple case of a single-item auction with two bidders, any mechanism can be forced to perform arbitrarily badly with respect to this measure. Instead, the second-lowest actual cost is the natural baseline for single-item auctions.

In order to analyze mechanisms for the "hiring a team" problem, this notion needs to be extended to set systems. One definition suggested and used previously in the literature [4, 27, 10] compares the payments $p(S)$ made to the winning feasible set S with the actual cost $c(S')$ of the cheapest feasible set S' disjoint from S; their ratio $p(S)/c(S')$ is called the *frugality ratio*, and measures the overpayment of the mechanism. However, this measure is not useful when there is no such disjoint set S' (and hence $c(S')$ may be assumed to be infinite). This may happen even when no element of S has a monopoly, as evidenced by the case of computing an MST on a cycle, or a shortest s-t path in the graph depicted in Figure 1 below:[2] Thus, for large classes of path auctions and MST auctions (let alone other "hiring a team" problems), up to this point, the question of how to define the problem of evaluating the frugality of truthful mechanisms has been unresolved.

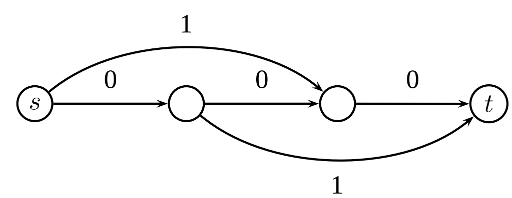

Figure 1. A graph with no monopoly and no solution disjoint from the cheapest one.

1.2 Our Results

As our first contribution, we present a new and natural definition of the frugality ratio of a mechanism. Our new definition is essentially based on the Nash Equilibrium of a non-truthful mechanism. We propose this as a benchmark against which to compare truthful mechanisms. The bound attempts to capture the minimum payments we could expect any mechanism, truthful or untruthful, to pay. In many cases, our new definition coincides with the one in [4, 27], but also extends it in a natural way to all monopoly-free collections of feasible sets.

[1] We assume here and in the sequel that there actually is a second-lowest bid. If there is only one bidder, or, more generally, if one agent is part of all feasible solutions, then that agent has a *monopoly*, and can dictate an arbitrary price. This case is not very interesting from the point of view of analysis, and not very pleasant from the point of view of the auctioneer. Hence, we focus only on *monopoly-free* instances.

[2] We thank Éva Tardos for suggesting this example.

With this new definition in hand, we proceed to investigate mechanisms for several subclasses of the "hiring a team" problem. First, we extend known results about MST auctions [5, 10], showing that the VCG mechanism has frugality ratio 1 for all two-connected (i.e., monopoly-free) graphs. The main result we prove here is that VCG has frugality ratio 1 if and only if the feasible sets are the maximal independent sets of a monopoly-free matroid.

We then turn to the study of shortest path auctions. Archer and Tardos [4], and Elkind et al. [8] have shown that if the graph consists of just two node-disjoint s-t paths of length $n/2$ each, then not only does VCG have a frugality ratio of $\Omega(n)$, but so does any truthful mechanism. As $O(n)$ is trivially an upper bound on the frugality ratio in all monopoly-free settings, this settles the worst-case frugality ratio of shortest-path auctions on worst-case graphs.

However, this worst case in no way applies to all graphs (see e.g. [18, 9]). The frugality ratio of a mechanism should therefore be bounded in terms of properties of the specific set system under consideration (in this case, the specific graph), rather than by the worst case over a class of set systems. We present a new, polynomial-time implementable mechanism, called the $\sqrt{\ }$-Mechanism, whose frugality ratio we show to be within a constant factor of optimal, for any given graph. In contrast, VCG's frugality ratio for some graphs can be as much as $\Omega(\sqrt{n})$ larger than that of our mechanism. Hence, VCG is in a sense *far* from the best possible for shortest path auctions on many graphs.

At the heart of the $\sqrt{\ }$-Mechanism for path auctions lies a simple mechanism for deciding which of two disjoint sets of different sizes should win. Similarly, the VCG mechanism for selecting an MST on a cycle can be seen as choosing $n-1$ out of n sets of size 1. A common generalization of both is to select r out of k sets (which may have different sizes). For such r-out-of-k systems, we present a mechanism whose frugality ratio is at most twice that of the optimum mechanism.

1.3 Related Work

Motivated by the need to deal with selfish users, in particular in network settings, there has been a large body of recent work at the intersection of game theory, economic theory and theoretical computer science (see, e.g., [20, 25]). For instance, the seminal paper of Nisan and Ronen [22], which introduced mechanism design to the theoretical computer science community, studied the tradeoffs between agents' incentives and computational complexity. The loss of efficiency in network games due to selfish user behavior has been studied in the contexts of the "price of anarchy" (see, e.g., [25, 26]), and the "price of stability" (see [2]).

The problem of hiring a team of agents in complex settings, at minimum total cost, has been shown to have many practical economic applications (see [9, 1, 23, 16, 21] for examples). In particular, the path auction problem we study here has been the subject of a significant amount of prior research. The traditional economics approach to payment minimization (or profit maximization) is to construct the optimal Bayesian auction given the prior distributions from which agents' private values are drawn. Indeed, path auctions and similar problems have been studied recently from the Bayesian perspective in [8, 7]. By way of contrast, we follow the approach pioneered by Archer, Tardos, Talwar and others [4, 27, 5], and study the problem of hiring a team from a worst-case perspective. As we have repeatedly seen in computer science, significant insight can be gained from an understanding of worst-case performance, and it enables an uninformed or only partially informed auctioneer to evaluate the trade-off between an auction tailored to assumptions about bidder valuations (which may or may not be correct) versus an auction designed to work as well as possible under unknown and worst-case market conditions.

2 Preliminaries

We formally define the class of "hiring a team" auctions we study. A *set system* $(E, \mathcal{F})$ is specified by a set E of n *elements*, each representing an *agent*, and a collection $\mathcal{F} \subseteq 2^E$ of *feasible sets*. We assume that $(E, \mathcal{F})$ is common knowledge to the auctioneer and all agents. A set system is *monopoly-free* if no element is in all feasible sets, i.e., if $\bigcap_{S \in \mathcal{F}} S = \emptyset$. Each agent $e \in E$ has a *cost* c_e, its true incurred cost if it is selected by the mechanism.[3] This value is private, i.e., known only to agent e. An *auction* consists of two steps:

1. Each agent submits a sealed bid b_e.
2. Based on the bids b_e, the auctioneer selects a feasible set $S \in \mathcal{F}$ as the winner, and computes a payment $p_e \geq b_e$ for each agent $e \in S$. We say that the agents $e \in S$ *win*, and all other agents *lose*.

Both the selection rule and the algorithm for computing payments are assumed to be common knowledge among the agents. We assume that the agents will exploit this knowledge to choose a bid maximizing their own *profit*, where an agent's profit is $p_e - c_e$ if the agent wins, and 0 otherwise. We also assume that agents do not collude.

A mechanism is *truthful* if, for any fixed vector b_{-e} of bids by all agents other than e, it is in agents e's best interest to bid $b_e = c_e$, i.e., agent e's profit is maximized by bidding $b_e = c_e$. Hence, we can use b_e and c_e interchangeably in discussing truthful mechanisms.

[3]For costs, bids, etc., we extend the notation by writing $c(S) = \sum_{e \in S} c_e$, etc.

It is well-understood which selection rules yield truthful mechanisms. In fact, the selection rule uniquely determines the payments for truthful mechanisms.

Theorem 1 *[3, 15]*

1. *A mechanism is truthful only if the selection rule is* monotone*: No losing agent can become a winner by raising his bid, given fixed bids by all other agents.*

2. *Given a monotone selection rule, there is a unique truthful mechanism with this selection rule. This mechanism pays each agent his* threshold bid*, i.e., the highest value he could have bid and still won (all other agents' bids being fixed).*

Theorem 1 will allow us to describe truthful mechanisms by only specifying a monotone selection rule, which implicitly defines the payments. We will also use the following immediate corollary.

Corollary 2 *If element e wins with costs c under a truthful mechanism $\mathcal{M}$, then e must also win with cost $c'_e \leq c_e$, when the costs of all other elements e' are unchanged, i.e. $c'_{e'} = c_{e'}$. Moreover, with costs c', the payment by $\mathcal{M}$ to e must be at least c_e.*

2.1 Nash Equilibria and Frugality Ratios

We define a new notion of the frugality ratio of a mechanism. Intuitively, the frugality ratio should capture the overpayment of a mechanism, with respect to a "natural" lower bound. We argue that a natural choice for this lower bound is essentially the minimum payment by a non-truthful mechanism; hence, the frugality ratio characterizes the cost of insisting on truthfulness.

Consider the mechanism $\mathcal{N}$ which, given the bids b_e, selects the cheapest feasible set with respect to these bids (using some tie-breaking rule), and pays each winning agent his bid. This mechanism is a *"first-price auction"* and is *not* truthful, and hence has no dominating strategies. In the absence of dominating strategies, it is common to study *Nash Equilibria*, bid vectors $\mathbf{b}$ such that no agent e, given the bids b_{-e} by all other agents, can increase his profit by changing his bid to something other than b_e. Hence, Nash Equilibria could be considered "natural outcomes" of the mechanism $\mathcal{N}$, and the resulting payments thus good candidate bounds. As we are interested in a lower bound, we wish to define the *cheapest Nash value* $\nu(\mathbf{c})$ to be the minimum payments by $\mathcal{N}$ over all of its Nash Equilibria.

Unluckily, as observed by Immorlica et al. [13], Nash Equilibria often do not exist for first-price auctions. We will discuss this issue in more detail below; however, we can still define the quantity $\nu(\mathbf{c})$ analytically based on the intuition we gained from the concept of Nash Equilibria.

Definition 3 *Let $(E, \mathcal{F})$ be a set system, and S the cheapest feasible set with respect to the true costs c_e (where ties are broken lexicographically). We define $\nu(\mathbf{c})$ to be the solution to the following optimization problem.*

Minimize $\sum_{e \in S} b_e$ *subject to*

(1) $b_e \geq c_e$ *for all e*

(2) $\sum_{e \in S \setminus T} b_e \leq \sum_{e \in T \setminus S} c_e$ *for all* $T \in \mathcal{F}$

(3) *For every $e \in S$, there is a $T_e \in \mathcal{F}$ such that $e \notin T_e$ and* $\sum_{e' \in S \setminus T_e} b_{e'} = \sum_{e \in T_e \setminus S} c_e$

The first constraint simply states that no agent is willing to incur a loss, and the second constraint ensures that the set S indeed has the lowest total bid among feasible sets. Finally, the third constraint states that for each element, there is a reason why it cannot raise its bid further to increase its profit. We say that the set T_e in the third constraint is *tight for* e or simply *tight*.

To illustrate this bound, consider the case of an s-t path auction in which there are k parallel paths. Then, $\nu(\mathbf{c})$ is precisely the cost of the second-cheapest path. Our intuition, derived from the notion of Nash Equilibria, is that the agents on the cheapest path will raise their bids until the sum of their bids equals the cost of the second-cheapest path, at which point they can no longer raise their bids. None of the other edges have an incentive to raise their bids (as they are losing either way), nor to lower their bids, as they would incur a negative profit. Thus, the new metric in this case is identical to that used in previous studies of path auctions. This coincidence extends to many set systems and cost vectors.

Unfortunately, the intuition that the quantity $\nu(\mathbf{c})$ from Definition 3 is the value of the true cheapest Nash is not quite correct (although we will sometimes abuse terminology and refer to $\nu(\mathbf{c})$ as the cheapest Nash value). As noted by Immorlica et al. [13], depending on the tie-breaking of the first price auction, there may not even be a Nash Equilibrium to a first-price auction. Moreover, even when the Nash Equilibrium does exist, the one yielding lowest total payments may have a winning set different from the actually cheapest set S. There are several technical ways to deal with these issues, including: (1) We could consider *strong ϵ-Nash equilibria*, in which there is no group of agents that can deviate in a way that improves the payoff of each member by at least ϵ. We could then take the limit for $\epsilon \to 0$. (2) We could consider strong Nash equilibria for a first price auction that uses oracle access to the true costs of the bidders to break ties. For a discussion and results about these and related issues in the context of obtaining upper bounds on the payments in first price auctions, see [13]. We do not elaborate on these issues here as they are incidental to our main goal of providing a well-motivated lower bound on the payments of any mechanism — instead, we use as our new metric the very intuitive value given by Definition 3.

This metric is defined in all monopoly-free set systems. As mentioned above, for path auctions with several node-disjoint s-t paths, $\nu(\mathbf{c})$ is equal to the cost of the second cheapest path. In the graph of Figure 1, the value of ν is 1, and occurs when the middle edge bids 1, and all other edges bid their true cost. For the spanning tree auction on a cycle of n nodes, where one edge has true cost 1 and the rest have cost 0, the value of $\nu(\mathbf{c})$ is $n-1$.

Using the quantity $\nu(\mathbf{c})$, we can now proceed to define the frugality ratio of a mechanism $\mathcal{M}$ for a set system $(E, \mathcal{F})$.

Definition 4 *Let $\mathcal{M}$ be a truthful mechanism for the set system $(E, \mathcal{F})$ and let $p_{\mathcal{M}}(\mathbf{c})$ denote the total payments of $\mathcal{M}$ when the actual costs are $\mathbf{c}$. Then, the* frugality ratio *of $\mathcal{M}$ is*

$$\phi_{\mathcal{M}} \quad = \quad \sup_{\mathbf{c}} \frac{p_{\mathcal{M}}(\mathbf{c})}{\nu(\mathbf{c})}.$$

The frugality ratio of the set system $(E, \mathcal{F})$ is $\Phi_{(E,\mathcal{F})} = \inf_{\mathcal{M}} \phi_{\mathcal{M}}$, where the infimum is taken over all truthful mechanisms $\mathcal{M}$ for the set system $(E, \mathcal{F})$. We call a mechanism $\mathcal{M}$ competitive *for a class of set systems $\{(E_1, \mathcal{F}_1), (E_2, \mathcal{F}_2), \ldots\}$ if $\phi_{\mathcal{M}}$ is within a constant factor of $\Phi_{(E_i, \mathcal{F}_i)}$ for all i.*

3 Spanning Trees and other Matroids

In this section, we generalize the results of Talwar [27, 5] and show that under the new definition of frugality ratio, the VCG mechanism has frugality ratio 1 if and only if the feasible sets are the maximal independent sets (bases) of a monopoly-free matroid. In particular, VCG has frugality ratio 1 for spanning tree auctions, even when there are no two disjoint spanning trees (e.g., when the graph is a cycle).

Interestingly, under the previous definition of frugality ratio, in which the payments were compared to the cost of the cheapest solution disjoint from the true cheapest solution, the set systems with frugality ratio 1, termed *frugoids*, were a strict superset of matroids [27]. With our new definition, this is no longer the case.

Theorem 5 *VCG has frugality ratio 1 for an auction on the set system $(E, \mathcal{F})$ if and only if the feasible sets $\mathcal{F}$ are the bases of a monopoly-free matroid.*

The proof is quite similar to the one in [27] for frugoids. It depends on the following well-known characterization of matroids [24]:

Proposition 6 *A collection $\mathcal{F}$ of sets forms the bases of a matroid if and only if for every two sets $S, T \in \mathcal{F}$, there is a bijection f between $S \setminus T$ and $T \setminus S$ such that $S \setminus \{e\} \cup \{f(e)\}$ is in $\mathcal{F}$ for every $e \in S \setminus T$.*

A central part of the proof will be relating the VCG payments to the bids b_e in Definition 3.

Proposition 7 *The VCG payment to each winning element e is at least the bid b_e from Definition 3.*

Proof. Let S be the winning set, and $e \in S$ arbitrary. VCG's payment to e is $p_e = c_e + c(T) - c(S)$, for some feasible set T with $e \notin T$. Because $b_{e'} \geq c_{e'}$ for all elements e', and $b(S) \leq c(T)$, we obtain $b_e \leq c(T) - b(S \setminus \{e\}) \leq c(T) - c(S \setminus \{e\}) = p_e$. ∎

Proof of Theorem 5. For the "if" direction, assume that the feasible sets are bases of a monopoly-free matroid. Fix an arbitrary cost vector $\mathbf{c}$, and let S be the winning set under the VCG mechanism. So S has minimum cost with respect to $\mathbf{c}$.

For each $e \in S$, let T_e be the set which is tight for e, i.e., $b(S \setminus T_e) = c(T_e \setminus S)$. By Proposition 6, there is a bijection f between $S \setminus T_e$ and $T_e \setminus S$, such that $S \setminus \{e'\} \cup \{f(e')\}$ is feasible for every $e' \in S \setminus T_e$, implying that $b_{e'} \leq c_{f(e')}$. Together with the above equality, this in fact implies that $b_{e'} = c_{f(e')}$ for all $e' \in S \setminus T_e$. VCG pays at most $c_{f(e')}$ to each element $e' \in S \setminus T_e$, so VCG has frugality ratio 1.

For the "only if" direction, assume that $(E, \mathcal{F})$ is not a matroid, and let S, T be two sets violating the condition of Proposition 6, with $|T| \leq |S|$. If $|T| < |S|$, then let $\mathbf{c}$ be such that all elements of S have cost 0, all elements of $T \setminus S$ have cost 1, and all elements outside $S \cup T$ have very large cost. Then, S is the cheapest set, and hence also the winning solution under VCG. Because T is also feasible, $b(S \setminus T) \leq c(T \setminus S) = |T \setminus S|$ in Definition 3. On the other hand, VCG has to pay each element $e \in S \setminus T$ at least 1, so its total payment to $S \setminus T$ is at least $|S \setminus T| > |T \setminus S|$ by assumption. For the elements in $S \cap T$, VCG pays at least b_e by Proposition 7, so VCG's total payments are strictly greater than $b(S)$.

When $|S| = |T|$, consider the bipartite graph with node sets $S \setminus T$ and $T \setminus S$, containing an edge (e, e') (for $e \in S, e' \in T$) if and only if $S \setminus \{e\} \cup \{e'\}$ is feasible. By the assumption and Hall's Theorem, there is a set $S' \subseteq S \setminus T$ such that its neighborhood $T' = \Gamma(S')$ in the bipartite graph has size $|T'| < |S'|$. Let $\mathbf{c}$ be the cost vector in which all elements of T' have cost 2, all elements of $T \setminus T'$ have cost 1, all elements of $S \setminus T$ have cost 0, and all elements outside $S \cup T$ have very large cost. Again, S is the cheapest set, and thus wins under VCG. Because T is feasible, the bids b of $S \setminus T$ in Definition 3 are at most

$$b(S \setminus T) \quad \leq \quad c(T \setminus S) \quad = \quad |T \setminus S| + |T'|.$$

On the other hand, any single element $e \in S'$ cannot be replaced by a single element of $T \setminus T'$, so VCG must pay each $e \in S'$ at least 2, and each $e \in S \setminus (S' \cup T)$ at least 1.

Hence, the total payments by VCG to $S \setminus T$ are at least

$$|S \setminus T| + |S'| \quad = \quad |T \setminus S| + |S'| \quad > \quad c(T \setminus S).$$

Using Proposition 7 for the elements in $S \cap T$ now completes the proof that VCG's total payments are strictly greater than $b(S)$. ■

4 r-out-of-k Systems

In this section, we develop a competitive mechanism with near-optimal frugality ratio for an interesting general class of symmetric set systems. The solution we develop for this case will help us obtain competitive mechanisms for path auctions in Section 5.

Definition 8 *A set system $(E, \mathcal{F})$ is an r-*out-of-*k-*system *if there exists a partition of E into k disjoint sets $S_1, \ldots, S_k$, such that every set $F \in \mathcal{F}$ contains exactly r out of these k sets of elements.*

Notice that these set systems generalize both spanning trees on a cycle (with all sets S_i having size 1 and $r = k-1$), and a shortest path chosen from among k vertex-disjoint paths (with $r = 1$).

Our mechanism for these set systems is based on the observation that if a large set is chosen as a winner, then *each* of its elements will have to be paid a certain amount (depending on the other sets' bids). Hence, to avoid overpayments, a mechanism should — within reason — give preference to smaller sets. Our mechanism achieves this by comparing not the costs of sets themselves, but the costs $c(S_i)$ weighted by coefficients γ_i which capture the relative size of S_i compared to other sets. The precise magnitude of the γ_i is chosen to balance the worst-case frugality ratio over all potential combinations of winning sets. Specifically, for a $(k-1)$-out-of-k-system, we define the γ_i as the solution to the system of equations

$$\alpha \quad = \quad \frac{1}{(k-1)\gamma_i} \cdot \sum_{j \neq i} \gamma_j |S_j|, \tag{1}$$

for $i = 1, \ldots, k$. We begin by proving that such γ_i actually exist.

Lemma 9 *There is a solution $\alpha, \gamma_1, \ldots, \gamma_k$ to the system of k equations (1).*

Proof. We write $s_i = |S_i|$. Multiplying the i^{th} equation with $(k-1)\gamma_i$, subtracting the i^{th} equation from the first one, and rearranging shows that we can equivalently write

$$\begin{array}{lcll} (k-1)\alpha\gamma_1 & = & \sum_{i=2}^{k} \gamma_i s_i, & \\ \gamma_i & = & \frac{s_1 + (k-1)\alpha}{s_i + (k-1)\alpha} \cdot \gamma_1 & \text{for all } i \geq 2. \end{array}$$

Substituting γ_i into the first equation, canceling out γ_1, and multiplying by the common denominator, we obtain the following k^{th} degree polynomial equation for α:

$$\begin{array}{l} (k-1)\alpha \cdot \prod_{j=2}^{k} (s_j + (k-1)\alpha) \\ - \sum_{j=2}^{k} s_j \cdot \prod_{i \neq j} (s_i + (k-1)\alpha) \quad = \quad 0. \end{array}$$

For $\alpha = 0$, the polynomial is negative, whereas for sufficiently large α, it is positive (because the coefficient of α^k is positive). Therefore, there must be a solution α, and hence a solution to the original system of equations (for instance fixing arbitrarily $\gamma_1 = 1$). ■

We now use these values in order to define our mechanism. To deal with a general r-out-of-k-system, we discard the most expensive $k - r - 1$ sets, and then consider only the problem for the r-out-of-$(r+1)$-system. Our full mechanism is thus:

> **r-out-of-k-system Mechanism:**
>
> **Input:** r-out-of-k-system $S_1, S_2, \ldots, S_k$, with bids b_e for each element in E.
>
> 1. Rename the sets S_i so that they are sorted by non-decreasing sums of bids $b(S_i)$.
>   (Discard all but the $r + 1$ cheapest sets.)
> 2. For $1 \leq i \leq r + 1$, let γ_i be defined by Equation (1) (with $k = r + 1$).
> 3. Let $\ell = \operatorname{argmax}_i \frac{b(S_i)}{\gamma_i}$.
>
> **Output:** $F = \cup_{i \neq \ell} S_i$

Theorem 10 *For any r between 1 and $k-1$, the r-out-of-k-system mechanism is truthful and competitive (i.e., achieves a frugality ratio within a constant factor of optimal).*

Proof. Since the selection rule is monotone, it follows from Theorem 1 that the r-out-of-k-system mechanism is truthful. Hence, we can assume that $b_e = c_e$ for all elements e.

The rest of the proof consists of two parts. We first show a lower bound of $\frac{\alpha}{2}$ on the frugality ratio of any truthful mechanism for a $(k-1)$-out-of-k-system. This lower bound extends straightforwardly to an r-out-of-k-system. To show competitiveness, we then verify that our mechanism has a frugality ratio of at most α.

Lemma 11 *Let $(E, \mathcal{F})$ be a $(k-1)$-out-of-k-system with set sizes $|S_1|, \ldots, |S_k|$, and let α be defined by the system of equations (1). Then $\frac{\alpha}{2}$ is a lower bound on the frugality ratio of any truthful mechanism for this set system.*

Proof. Let $\mathcal{M}$ be any truthful mechanism for $(E, \mathcal{F})$. We define a directed graph G with node weights as follows: The nodes are the elements $e \in E$, and element e has weight $w(e) = \gamma_{i(e)}$, where $i(e)$ is the unique index such that $e \in S_{i(e)}$, and the γ_i are defined by the system of equations (1). The graph contains an arc (e, e') if node e' wins under the mechanism $\mathcal{M}$ when element e bids $\gamma_{i(e)}$, element e' bids $\gamma_{i(e')}$, and all other nodes bid 0. Notice that by the definition of the set system, the graph contains at least one of the two arcs (e, e') or (e', e) whenever $i(e) \neq i(e')$. Thus, for any node e, $\sum_{e':e\to e'} w(e') + \sum_{e':e'\to e} w(e') \geq \sum_{j\neq i(e)} \gamma_j |S_j|$.

Any node-weighted directed graph contains a node v such that $\sum_{u:v\to u} w(u) \geq \sum_{u:u\to v} w(u)$ so there is a node e (and set $S_i \ni e$) satisfying $\sum_{e':e\to e'} w(e') \geq \frac{1}{2}\sum_{j\neq i} \gamma_j |S_j|$.

Now consider what happens when the mechanism $\mathcal{M}$ is run with element e's cost being $c_e = \gamma_{i(e)}$, and all other costs being 0. By definition of the graph G and Corollary 2, all elements e' with an incoming arc from e win, and their threshold bid is at least their weight. Hence, the total payments by $\mathcal{M}$ are at least $\frac{1}{2}\sum_{j\neq i} \gamma_j |S_j|$. For this input $\mathbf{c}$, we have $\nu(\mathbf{c}) = (k-1)\gamma_i$, so the frugality ratio of $\mathcal{M}$ is at least $\frac{\frac{1}{2}\sum_{j\neq i}\gamma_j|S_j|}{(k-1)\gamma_i} = \frac{\alpha}{2}$. ■

Let e be a winning element in our r-out-of-k-system mechanism. Assume e is in set S_i. The payment to e is its threshold bid, i.e., the largest value e could have bid and still won. If e's bid were so large that the set S_i (instead of S_ℓ) were to become the argmax in step 3 of our mechanism, then e would clearly lose, so we can upper-bound the payment to e by $\frac{\gamma_i}{\gamma_\ell} c(S_\ell)$. Thus the total payment is upper-bounded by $\sum_{1\leq i\leq r+1, i\neq \ell} \frac{\gamma_i |S_i|}{\gamma_\ell} c(S_\ell)$.

On the other hand, we have $\nu(\mathbf{c}) = r \cdot c(S_{r+1})$. Since $c(S_\ell) \leq c(S_{r+1})$, the frugality ratio is upper-bounded by $\frac{1}{r\gamma_\ell} \cdot \sum_{1\leq i\leq r+1, i\neq\ell} \gamma_i |S_i| = \alpha$.

Finally, we can show that α is a lower bound on the frugality ratio $\Phi_{(E,\mathcal{F})}$ for this set system, by considering only bid vectors where all elements of S_j for $j > r+1$ have very large (essentially infinite) cost, so that no frugal mechanism will ever include them. Then, we can apply Lemma 11 to obtain the desired lower bound, completing the proof. ■

4.1 A tighter lower bound for 1-out-of-2-systems

The lower bound from Lemma 11, when applied to 1-out-of-2 systems, shows that any truthful mechanism for selecting exactly one of two disjoint sets S_1, S_2 must have frugality ratio at least $\frac{1}{2}\sqrt{|S_1||S_2|}$, which matches the upper bound to within a factor of 2. We now show that the lower bound can be tightened somewhat.

Lemma 12 *Any truthful mechanism $\mathcal{M}$ for a 1-out-of-2 system with sets S_1 and S_2 has frugality ratio at least $\frac{1}{\sqrt{2}}\sqrt{|S_1||S_2|}$.*

We conjecture that any truthful mechanism must have frugality ratio at least $\sqrt{|S_1||S_2|}$, i.e., that our mechanism is optimal. It can be shown to be optimal within the smaller class of *Min-Function mechanisms*, defined by Archer and Tardos [4].

Proof. Let $s_1 = |S_1|, s_2 = |S_2|$, and let $e_1, \ldots, e_{s_1}$ and $e'_1, \ldots, e'_{s_2}$ denote the elements of S_1 resp. S_2. We let a_{ij} denote the supremum of all values a such that if element e_i bids a, element e'_j bids 1, and all others bid 0, then set S_1 wins under the mechanism $\mathcal{M}$.

If element e'_j has cost 1, and all other elements have cost 0, then by Corollary 2, the set S_1 wins, and each element e_i is paid at least $a_{ij} - \epsilon$, for arbitrarily small ϵ. Hence, the total payments, and thus also the frugality ratio, are lower bounded by $\sum_i a_{ij}$ (omitting the ϵ).

Conversely, fix an element e_i, and consider a non-decreasing sorting A_{ip} of the a_{ij}, i.e. A_{ip} is the p^{th} smallest one among the a_{ij}. Then, if element e_i has cost $A_{ip} + \epsilon$, and all other elements have cost 0, the set S_2 wins, and each element e'_j with $a_{ij} \leq A_{ip}$ is paid at least 1. Thus, the total payment is at least p, whereas $\nu(\mathbf{c}) = A_{ip} + \epsilon$, so the frugality ratio is lower-bounded by $\frac{p}{A_{ip}}$ for all i, p (again omitting ϵ). The frugality ratio $\phi_{\mathcal{M}}$ therefore satisfies the following inequalities:

$$\begin{array}{lcll} \sum_i a_{ij} & \leq & \phi_{\mathcal{M}} & \text{for all } j \\ A_{ip} & \geq & \frac{p}{\phi_{\mathcal{M}}} & \text{for all } i, p. \end{array}$$

As each a_{ij} appears as exactly one A_{ip}, their sums are equal, i.e.,

$$\frac{s_1}{\phi_{\mathcal{M}}} \cdot \frac{s_2(s_2+1)}{2} \quad \leq \quad \sum_{i,p} A_{ip} \quad = \quad \sum_{i,j} a_{ij} \quad \leq \quad s_2 \phi_{\mathcal{M}}.$$

Solving for $\phi_{\mathcal{M}}$ yields that $\phi_{\mathcal{M}}^2 \geq \frac{s_1(s_2+1)}{2}$, i.e.,

$$\phi_{\mathcal{M}} \quad \geq \quad \sqrt{s_1(s_2+1)/2} \quad \geq \quad \sqrt{s_1 s_2 / 2}.$$

■

5 Path Auctions Revisited

The mechanism for the 1-out-of-2-system can be used as a building block in order to get a near-optimal mechanism for path auctions. Recall that for the case $k = 2$, the solution to the system of equations (1) is particularly simple. For sets S_1, S_2 of sizes s_1, s_2, the solution is $\alpha = \sqrt{s_1 s_2}$, with $\gamma_1 = \frac{1}{\sqrt{s_1}}$ and $\gamma_2 = \frac{1}{\sqrt{s_2}}$.

Thus, for a directed graph consisting of two node-disjoint s-t paths P_1, P_2 of lengths s_1, s_2, Theorem 10 implies that the mechanism which selects the path P_1 iff $\sqrt{s_1} \cdot c(P_1) \leq \sqrt{s_2} \cdot c(P_2)$ has frugality ratio $\sqrt{s_1 s_2}$, and by Lemma 12, is within a factor of $\sqrt{2}$ of optimal. We call this mechanism the $\sqrt{\ }$*-Mechanism*. In contrast, the VCG mechanism has frugality ratio $\max(s_1, s_2)$ for this graph, and thus can perform by a factor of $\sqrt{n}$ worse in the case when $s_1 = 1$ and $s_2 = n$. (When $s_1 = s_2$, VCG and the $\sqrt{\ }$-Mechanism are identical. The deviation in behavior increases with the ratio $\frac{s_1}{s_2}$.)

The following generalized version of this mechanism is competitive for any monopoly-free graph.

$\sqrt{\ }$-Mechanism for $s - t$ Paths:

Input: A directed graph $G = (V, E)$ with bids b_e on the edges.

1. Find two edge-disjoint paths P, P' minimizing $b(P) + b(P')$.
   (Ignore the rest of the graph G from now on.)
2. Let $s = v_1, v_2, v_3, \ldots, v_{k+1} = t$ be the vertices that P, P' have in common, in the order in which they appear in P and P'. Let P_i (resp. P'_i) be the subpath of P (resp. P') from v_i to v_{i+1}.
3. For each i, include P_i in the solution iff $\sqrt{|P_i|} \cdot c(P_i) \leq \sqrt{|P'_i|} \cdot c(P'_i)$;
   otherwise, include P'_i.

Output: An s-t path consisting of the winners from each of the (P_i, P'_i) pairs.

Theorem 13 *The $\sqrt{\ }$-Mechanism for s-t paths is truthful, competitive (i.e., has frugality ratio within a constant factor of the optimal frugality ratio), and polynomial time implementable.*

Proof. Since the path selection rule is monotone, it follows from Theorem 1 that the $\sqrt{\ }$-Mechanism for s-t paths is truthful. Hence, we can assume that $b_e = c_e$ for all edges e.

The payments made to the winner among P_i, P'_i are at most $\sqrt{|P_i||P'_i|} \cdot \max(c(P_i), c(P'_i))$, and hence the total payments are at most $\sum_{i=1}^{k} \sqrt{|P_i||P'_i|} \cdot \max(c(P_i), c(P'_i))$. By Lemma 14 below, we have $\nu(\mathbf{c}) \geq \frac{1}{2}(c(P) + c(P')) \geq \frac{1}{2} \sum_{i=1}^{k} \max(c(P_i), c(P'_i))$. Hence, the $\sqrt{\ }$-Mechanism has frugality ratio at most $2 \max_i \sqrt{|P_i||P'_i|}$.

Let i be the index maximizing $|P_i||P'_i|$. By considering inputs with very large costs for all edges e outside $P \cup P'$ and costs 0 for all P_j, P'_j with $j \neq i$, we can apply Lemma 12 to the sets P_i and P'_i, obtaining a lower bound of $\frac{1}{\sqrt{2}} \sqrt{|P_i||P'_i|}$ for the frugality ratio of any truthful mechanism. Hence, the $\sqrt{\ }$-Mechanism is within a factor of $2\sqrt{2}$ of optimal for all graphs.

Finally, to see that the mechanism is polynomial time implementable, we observe that step 1 can be implemented in polynomial time by finding a minimum cost circulation where s has a demand of -2, t has a demand of 2 and all other nodes have demand 0, and that in step 3, any good approximation to the square root of the path lengths will change the frugality ratio of the mechanism by a corresponding (insignificant) amount. ■

To complete the proof of the theorem, it remains to establish the following lemma, bounding $\nu(\mathbf{c})$ from below.

Lemma 14 *Let P, P' be two edge-disjoint s-t paths minimizing $c(P) + c(P')$. Then, $\nu(\mathbf{c}) \geq \frac{1}{2}(c(P) + c(P'))$.*

Proof. Fix a vector $\mathbf{b}$ of bids from Definition 3 of $\nu(\mathbf{c})$. Let x be minimal such that there are at least two edge-disjoint s-t paths Q, Q', with costs at most x each. We will show that $x \leq \nu(\mathbf{c})$, which implies the lemma, as Q, Q' are candidate paths for P, P', and hence $\frac{1}{2}(c(P) + c(P')) \leq x$.

Assume for contradiction that $x > \nu(\mathbf{c})$. Then, the subgraph consisting only of paths with true costs at most $\nu(\mathbf{c})$ is not two-connected, and must have a cut-edge e. All s-t paths P minimizing $b(P)$ must have $\nu(\mathbf{c}) = b(P) \geq c(P)$, and hence must contain e. But then, there is no tight set for e, contradicting the third (tightness) constraint in Definition 3. ■

Lemma 14 only gives a lower bound on the value of $\nu(\mathbf{c})$, but no precise characterization. It is an interesting open question whether $\nu(\mathbf{c})$ can be characterized in terms of properties of the graph G and the costs $\mathbf{c}$, and whether it can be computed in polynomial time. The lower bound in Lemma 14 is off by at most a factor of 2, as it can be shown that the payments in *any* Nash Equilibrium are at most $c(P) + c(P') - c(\hat{P})$, where P, P' are two edge-disjoint s-t paths minimizing $c(P) + c(P')$, and $\hat{P}$ is a cheapest s-t path.

6 Conclusions and Open Problems

In this paper, we studied the problem of hiring a set of agents to perform a complex task, and the design and analysis of truthful mechanisms for interacting with selfish agents in this setting. We proposed a natural measure of the frugality of a mechanism, and revisited several known results under this new definition.

We propose studying mechanisms not in their worst-case behavior over large classes of set systems, but rather by analyzing them for individual set systems. Such a more detailed analysis exhibits a rich structure among mechanisms. In particular, we showed that the VCG mechanism is far from optimal for many instances of shortest-path or r-out-of-k auctions, and presented a different truthful mechanism that is always within a constant factor of optimal for these classes.

Our work should be considered as a first step toward understanding auctions in a worst-case setting for arbitrary set systems. The ultimate goal is to give natural, efficiently computable when possible, mechanisms that achieve a frugality ratio that is within a constant factor of optimum on every input. This problem is likely to be quite difficult, as there will be very complex interactions between agents. In order to extend our results to other set systems (and because it is of interest in its own right), it appears necessary to gain a deeper understanding of $\nu(\mathbf{c})$ for these set systems. Indeed, while we can pin its value down for shortest path auctions to within a factor of 2, we have no precise characterization; we do not even know whether it can be computed in polynomial time. For several problems, it seems that the main obstacle in proving competitiveness for natural extensions of the $\sqrt{}$-Mechanism is a lack of characterization of the value of $\nu(\mathbf{c})$.

The converse of designing new mechanisms is to ask for which set systems known mechanisms — say, VCG — are competitive. Notice that this would not require the frugality ratio to be 1 (as Theorem 5 characterizes those set systems); rather, it requires the frugality ratio to be within a constant factor of that of any other mechanism.

As we trace out the boundaries between the (traditional economics) Bayesian analysis for known cost distributions, and the (computer science) worst-case analysis, we may want to study the design of mechanisms given partial information: can our mechanisms be modified in a natural and principled way to take advantage of information such as an upper or lower bound on the cost of certain edges? Further natural extensions would consider the case where a single agent may own multiple elements (but obviously no monopoly), or where the auctioneer is constrained by a target budget B. In the latter case, the goal would be to purchase a feasible solution at an expense of at most B whenever there exists a feasible solution of cost at most αB for some $\alpha < 1$. Our goal would then be to find a mechanism achieving large α. For the case of s-t path auctions, for instance, it may be possible to adapt cost-sharing techniques [19] to achieve $\alpha = \Omega(\frac{1}{\log k})$, where k is the length of the cheapest path; on the other hand, we conjecture that it is not possible to achieve larger values of α.

Finally, as we mentioned above, the value of $\nu(\mathbf{c})$ should be investigated further. In particular, while it is a "natural" candidate for a lower bound on a non-truthful mechanism's payments, it would be desirable to come up with a more rigorous proof for this intuition. More importantly, perhaps, an extension of this concept to combinatorial auctions may help in providing a similar appropriate benchmark for the design of truthful, profit-maximizing combinatorial auctions.

Acknowledgments

We would like to thank Ning Chen, Edith Elkind, Jason Hartline, David Karger, Nati Linial, Evdokia Nikolova, Kunal Talwar, and Éva Tardos for useful discussions.

References

[1] E. Anderson, F. Kelly, and R. Steinberg. A contract and balancing mechanism for sharing capacity in a communication network. To appear.

[2] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th IEEE Symp. on Foundations of Computer Science*, 2004.

[3] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science*, 2001.

[4] A. Archer and E. Tardos. Frugal path mechanisms. In *Proc. 13th ACM Symp. on Discrete Algorithms*, pages 991–999. ACM/SIAM, 2002.

[5] S. Bikhchandani, S. de Vries, J. Schummer, and R. Vohra. Linear programming and vickrey auctions. *IMA Volume in Mathematics and its Applications, Mathematics of the Internet: E-auction and Markets*, 127:75–116, 2001.

[6] E. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[7] A. Czumaj and A. Ronen. On the expected payment of mechanisms for task allocation. In *Proc. 23rd ACM Symp. on Principles of Distributed Computing*. ACM Press, 2004.

[8] E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *Proc. 15th ACM Symp. on Discrete Algorithms*. ACM/SIAM, 2004.

[9] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker. A BGP-based mechanism for lowest-cost routing. In *Proc. 21st ACM Symp. on Principles of Distributed Computing*, 2002.

[10] R. Garg, V. Kumar, A. Rudra, and A. Verma. Coalitional games on graphs: core structures, substitutes and frugality. Technical Report TR-02-60, UTCS, 2002.

[11] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. Technical Report STAR-TR-99.09.01, STAR Laboratory, InterTrust Tech. Corp., Santa Clara, CA, 1999. Available at URL http://www.star-lab.com/tr/tr-99-01.html.

[12] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[13] N. Immorlica, D. Karger, E. Nikolova, and R. Sami. First-price path auctions. In *Proc. 7th ACM Conf. on Electronic Commerce*, 2005.

[14] P. Klemperer. Auction theory: A guide to the literature. *Journal of Economic Surveys*, 13:227–286, 1999.

[15] V. Krishna. *Auction Theory*. Academic Press, 2002.

[16] A. Lazar, A. Orda, and D. Pendarakis. Virtual path bandwidth allocation in multi-user networks. In *Proc. 14th IEEE INFOCOM Conference*, 1995.

[17] A. Mas-Collel, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.

[18] M. Mihail, C. Papadimitriou, and A. Saberi. On certain connectivity properties of the internet topology. In *Proc. 35th ACM Symp. on Theory of Computing*. ACM Press, 2003.

[19] H. Moulin and S. Shenker. Strategyproof sharing of submodular costs. *Journal of Economic Theory*, to appear.

[20] N. Nisan. Algorithms for selfish agents: Mechanism design for distributed computation. In *Proc. 17th Annual Symp. on Theoretical Aspects of Computer Science*, 1999.

[21] N. Nisan, S. London, O. Regev, and N. Carmiel. Globally distributed computation over the internet - the popcorn project. In *Proc. of the 18th International Conference on Distributed Computing Systems*, 1998.

[22] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proc. 31st ACM Symp. on Theory of Computing*, pages 129–140. ACM Press, 1999.

[23] D. O'Neill, D. Julian, M. Chiang, and S. Boyd. QoS and fairness constrained convex optimization of resource allocation for wireless, cellular and ad hoc networks. In *Proc. 21st IEEE INFOCOM Conference*, 2002.

[24] J. Oxley. *Matroid Theory*. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1992.

[25] C. Papadimitriou. Algorithms, games and the internet. In *Proc. 33rd ACM Symp. on Theory of Computing*, pages 749–752. ACM Press, 2001.

[26] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 2002.

[27] K. Talwar. The price of truth: Frugality in truthful mechanisms. In *Proc. 21st Annual Symp. on Theoretical Aspects of Computer Science*, 2003.

[28] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance*, 16:8–37, 1961.