

The Price of Truth: Frugality in Truthful Mechanisms

Kunal Talwar *

University of California, Berkeley CA 94720, USA,
kunal@cs.berkeley.edu

Abstract. The celebrated Vickrey-Clarke-Grove (VCG) mechanism induces selfish agents to behave truthfully by paying them a premium. In the process, it may end up paying more than the actual cost to the agents. For the minimum spanning tree problem, if the market is “competitive”, one can show that VCG never pays too much. On the other hand, for the shortest s - t path problem, Archer and Tardos [5] showed that VCG can overpay by a factor of $\Omega(n)$. A natural question that arises then is: For what problems does VCG overpay by a lot? We quantify this notion of overpayment, and show that the class of instances for which VCG never overpays is a natural generalization of matroids, that we call *frugoids*. We then give some sufficient conditions to upper bound and lower bound the overpayment in other cases, and apply these to several important combinatorial problems. We also relate the overpayment in an suitable model to the *locality ratio* of a natural local search procedure.

Classification: Current challenges, Mechanism design.

1 Introduction

Many problems require the co-operation of multiple participants, e.g. several autonomous systems participate to route packets in the Internet. Often these participants or *agents* have their own selfish motives, which may conflict with social welfare goals. In particular, it may be in an agent’s interest to misrepresent her utilities/costs. The field of *mechanism design* deals with the design of protocols which ensure that the designer’s goals are achieved by incentivizing selfish agents to be truthful. A *mechanism* is a protocol that takes the announced preferences of a set of agents and returns an *outcome*. A mechanism is *truthful* or *strategy proof* if for every agent, it is most beneficial to reveal her true preferences.

Consider, for example the problem of choosing one of several contractors for a particular task. Each contractor bids an amount representing her cost for performing the task. If we were to choose the lowest bidder, and pay her what her bid was, it might be in her interest to bid higher than her true cost, and make a large profit. We may, on the other hand, use the *VCG mechanism* [24, 8,

* Supported in part by NSF grants CCR-0105533 and CCR-9820897.

11] which in this case, selects the lowest bidder, and pays her an amount equal to the second lowest bid. It can be shown that in this case, it is in the best interest of every contractor to reveal her true cost (assuming that agents don't collude). Moreover, the task get completed at the lowest possible cost. The VCG mechanism is strategy proof, and minimizes the *social cost*. On the other hand, the cost to the mechanism itself may be large. In this case, for example, while the task gets completed by the most efficient contractor, the amount the mechanism has to pay to extract the truth is *more* than the true cost. In this paper, we address the question: How much more?

We look at a more general setting where some task can be accomplished by hiring a team of agents. Each agent performs a fixed service and incurs a fixed cost for performing that service. There are some given teams of agents, such that any one team can accomplish the complete task. The cost of each agent is known only to the agent herself, and agents are selfish. While the VCG mechanism selects a team that performs the task with minimum cost to itself, the amount that the mechanism pays to the agents may be high. We are interested in characterizing problems where this payment is *not too large*. Archer and Tardos [5] showed that in general, this payment can be $\Omega(n)$ times the real cost, even if there is sufficient competition.

It is not immediately clear what this payment should be compared to. Ideally, we would like to say that we never pay too much more than the real cost of the optimal. However, in a monopolistic situation, VCG can do pretty badly. So we impose the condition that the market is sufficiently competitive (a notion elucidated later), and we would be satisfied if our mechanism did well under these constraints. Qualitatively, we want to say that a mechanism is *frugal* for an instance if, in the presence of competition, the amount that the mechanism pays is not too much more than the true cost. Concretely let opt be the most cost-effective team for the task. We compare the amount that the mechanism pays to the agents in opt to the cost of the best rival solution opt' , i.e. the best solution to the instance that does not use any agent in opt . If it is the case that opt and opt' have exactly the same cost, it can be shown that VCG does not over-pay, i.e. it pays exactly the cost of opt . However, if opt' is a little costlier than opt , the performance may degrade rapidly. We define the *frugality ratio* of VCG on an instance to be the worst possible ratio of the payment to the cost of opt' .¹ This definition turns out to be equivalent to the *agents are substitutes* condition defined independently by [7] in a different context. We discuss the implications of this in section 5.

Under this definition of frugality, we characterize exactly the class of problems with frugality ratio 1. This class is a natural generalization of the class of matroids, which we call *frugoids*. We also give more general upper and lower bounds on the frugality in terms of some parameters of the problem. Using these, we classify several interesting problems by their frugality ratio. We also note a very interesting connection between the notion of frugality, and the *locality gap*

¹ We show in section 3.1 how this definition of frugality relates well to the notion of overpayment in the presence of competition.

of a natural local search procedure for the facility location problem; we discuss this in section 4.

Related Work

The problem of mechanism design has classically been a part of game theory and economics (see, e.g. [21], [18]). In the past few years, there has been a lot of work the border of computer science, economics and game theory (see [22] for a survey). Nisan and Ronen [19] applied the mechanism design framework to some optimization problems in computer science. Computational issues in such mechanisms have also been considered in various scenarios (see, e.g. [20], [10], [15], [7], [6]). Work has also been done in studying efficiently computable mechanisms other than VCG for specific problems (e.g. [16], [4], [17], [3]).

The issue of frugality is raised in [4] who look at a scheduling problem and compare the payment to the actual cost incurred by the machines under competitiveness assumptions. In [5, 9], the authors show that for the shortest path problem, the overpayment is large even in the presence of competition. On the positive side, [6] show that when the underlying optimization problem is a matroid, the frugality ratio is 1. [12] relates frugality to the core of a cooperative game.

The rest of the paper is organized as follows. In section 2, we define the framework formally. We state and prove our results on frugality in section 3. In section 4, we apply the results to several combinatorial problems, and in section 5, we consider some interesting implications. We conclude in section 6 with some open problems.

2 Definitions and Notation

We first define formally the problem for which we analyze the performance of VCG. We are given a set E of elements and a family $\mathcal{F} \subseteq 2^E$ of *feasible* subsets of E . We shall call (E, \mathcal{F}) a *set system*. A set system (E, \mathcal{F}) is *upwards closed* if for every $S \in \mathcal{F}$ and every superset $T : S \subseteq T \subseteq E$, it is the case that $T \in \mathcal{F}$. We also have a non-negative cost function $\hat{c} : E \rightarrow \mathbb{R}_0$. For a set $S \subseteq E$, let $\hat{c}(S) = \sum_{e \in S} \hat{c}(e)$. The goal is to find a feasible set of minimum total cost. Clearly for the minimization problem, there is no loss of generality in assuming that (E, \mathcal{F}) is upward closed. (E, \mathcal{F}) is therefore fully defined by the minimal sets in \mathcal{F} . We shall henceforth represent a set system by its minimal feasible sets (*bases*) $\mathcal{B} = \{S \in \mathcal{F} : \text{no proper subset of } S \text{ is in } \mathcal{F}\}$. Note that the family \mathcal{B} may be exponential in size and given implicitly. Several problems such as minimum spanning tree, shortest path, etc. can be put in this framework.

Now suppose that each element $e \in E$ is owned by a selfish agent, and she alone knows the true cost $c(e)$ of this element. A *mechanism* is denoted as $m = (\mathcal{A}, p)$ where

- \mathcal{A} is an allocation algorithm that takes as input the instance (E, \mathcal{B}) and the revealed costs $c(e)$ for each agent e , and outputs a set $S \in \mathcal{F}$.

- p is a payment scheme that defines the payment $p(e)$ to be made to each agent.

For the VCG mechanism, \mathcal{A} selects the minimum cost feasible set. The payment scheme p is as follows. Let $OPT(c)$ be the cost of an optimal solution under cost function c , i.e. $OPT(c) = \min_{S \in \mathcal{B}} c(S)$ and let $\text{opt}(c)$ denote the set achieving this minimum². Let $c[e \mapsto x]$ denote the cost function c' which is equal to c everywhere except at e , where it takes the value x . Then

$$p(e, c) = \begin{cases} 0 & \text{if } e \notin \text{opt}(c) \\ OPT(c[e \mapsto \infty]) - OPT(c[e \mapsto 0]) & \text{if } e \in \text{opt}(c) \end{cases}$$

It can be shown that the mechanism is *truthful* (see for example [18]). Hence assuming that agents are rational, $c(e) = \hat{c}(e)$ for all agents. We denote by $OPT'(c)$ the cost of the cheapest feasible set disjoint from $\text{opt}(c)$, i.e. $OPT'(c) = \min_{S \in \mathcal{B}, S \cap \text{opt}(c) = \emptyset} \{c(S)\}$. Note that such a set may not always exist, in which case this minimum is infinite. If $OPT'(c)$ is finite, let $\text{opt}'(c)$ denote the set achieving the minimum. For any set $S \subseteq E$, let $p(S, c) = \sum_{e \in S} p(e, c)$. For an instance (E, \mathcal{B}, c) , such that $OPT'(c) > 0$, define the *frugality ratio* $\phi(E, \mathcal{B}, c) = p(\text{opt}(c), c) / OPT'(c)$. (If $OPT'(c) = 0$, in which case $p(\text{opt}(c), c) = 0$ as well, we shall define the frugality ratio to be 1.) For a set system (E, \mathcal{B}) , define the frugality ratio of the the set system $\phi(E, \mathcal{B}) = \sup_c \phi(E, \mathcal{B}, c)$.

An alternative definition of frugality, what we call the *marginal frugality*, is the ratio of VCG overpayment to the difference in the costs of opt and opt' . Formally, $\phi'(E, \mathcal{B}) = \sup_c (p(\text{opt}(c), c) - OPT(c)) / (OPT'(c) - OPT(c))$. In the next section, we show how these two definitions are related. The following theorem shows that the marginal frugality of a set system bounds the rate of change of the payment with respect to the value of OPT' . Thus the marginal frugality behaves like the derivative of the payment with respect to $OPT'(c)$. We omit the proofs of the following from this extended abstract.

Theorem 1. *Let (E, \mathcal{B}) be a set system and let c and c' be cost functions such that $\text{opt}(c) = \text{opt}(c') = A$, $\text{opt}'(c) = \text{opt}'(c') = B$, and $c(e) = c'(e)$ for all $e \notin B$. Then*

$$p(A, c') - p(A, c) \leq \phi'(E, \mathcal{B}) \cdot (c'(B) - c(B))$$

Corollary 1.

$$\phi'(E, \mathcal{B}) = \lim_{\epsilon \rightarrow 0^+} \sup (p(A, c') - p(A, c)) / \epsilon$$

where the sup is over all c, c' satisfying the conditions in theorem 1 such that $c'(B) - c(B) = \epsilon$.

Thus, the “slope” of the payment function with respect to OPT' is always less than ϕ' , and is equal to ϕ' at $OPT' = OPT$.

² We assume that ties are broken in a particular way, say $\text{opt}(c)$ is the lexicographically smallest set $S \in \mathcal{B}$ such that $c(S) = OPT(c)$, so that $\text{opt}(c)$ is well defined.

3 Frugality

3.1 Canonical Cost functions

Call a cost function c *canonical* if the following hold:

- $c(e) = 0$ for all $e \in \text{opt}(c)$.
- $c(e) = \infty$ for all $e \notin \text{opt}(c) \cup \text{opt}'(c)$.

The following lemma shows that for every set system (E, \mathcal{B}) , and every cost function c , there exists a cost function c' which is canonical, and has a higher frugality ratio. Hence the definition of frugality ratio can be modified so as to consider only canonical cost functions.

Lemma 1. *Let (E, \mathcal{B}) be a set system, and c be a non negative cost function on E . Then there is a cost function c' such that the following hold:*

- c' is canonical.
- $\phi(E, \mathcal{B}, c') \geq \phi(E, \mathcal{B}, c)$

We omit the proof from this extended abstract.

It is easy to see that VCG payments satisfy the following properties.

Property 1. For every cost function c such that $OPT(c) = OPT'(c)$, $p(\text{opt}(c), c) = OPT(c)$.

Property 2. For every cost function c and every constant $\alpha > 0$, $p(\text{opt}(\alpha c), \alpha c) = \alpha p(\text{opt}(c), c)$.

The following lemma shows that under some constraints, VCG payments are super-additive.

Lemma 2. *Let c_1 be an arbitrary cost function and c_2 be a canonical cost function such that $\text{opt}(c_1) = \text{opt}(c_2)$ and $\text{opt}'(c_1) = \text{opt}'(c_2)$. Let $c = c_1 + c_2$. Then*

$$p(\text{opt}(c), c) \geq p(\text{opt}(c_1), c_1) + p(\text{opt}(c_2), c_2)$$

Proof. First note that $\text{opt}(c) = \text{opt}(c_1) = \text{opt}(c_2)$ and $\text{opt}'(c) = \text{opt}'(c_1) = \text{opt}'(c_2)$. Now let e be any element in $\text{opt}(c)$. Then

$$\begin{aligned} p(e, c) &= OPT(c[e \mapsto \infty]) - OPT(c[e \mapsto 0]) \\ &\geq c_1(\text{opt}(c[e \mapsto \infty])) + c_2(\text{opt}(c[e \mapsto \infty])) - c_1(\text{opt}(c)) - c_2(\text{opt}(c)) && \text{[by definition]} \\ &\geq c_1(\text{opt}(c_1[e \mapsto \infty])) + c_2(\text{opt}(c_2[e \mapsto \infty])) - c_1(\text{opt}(c)) - c_2(\text{opt}(c)) && \text{[Since } c[e \mapsto 0](\text{opt}(c)) \leq c(\text{opt}(c))\text{]} \\ &\geq c_1(\text{opt}(c_1[e \mapsto \infty])) + c_2(\text{opt}(c_2[e \mapsto \infty])) - c_1(\text{opt}(c)) && \text{[Since } c(\text{opt}(c)) \leq c(\text{opt}(c')), c_2 \text{ is canonical]} \\ &= (c_1(\text{opt}(c_1[e \mapsto \infty])) - c_1(\text{opt}(c_1[e \mapsto 0]))) + c_2(\text{opt}(c_2[e \mapsto \infty])) && \text{[Since } \text{opt}(c_1[e \mapsto 0]) = \text{opt}(c)\text{]} \\ &= p(e, c_1) + p(e, c_2) && \text{[} c_2 \text{ is canonical]} \end{aligned}$$

Summing over all e in $\text{opt}(c)$, we get the desired result.

Lemmas 1 and 2 imply:

Corollary 2. *Let (E, \mathcal{B}) be a set system. Then for any $\epsilon > 0$, there is a cost function c such that $OPT(c) = 1$, $OPT'(c) = (1 + \epsilon)$ and $p(\text{opt}(c), c) \geq 1 + \phi(E, \mathcal{B})\epsilon$.*

Proof. (Sketch) Let c be a canonical cost function achieving ³ a frugality ratio of $\phi(E, \mathcal{B})$. Let $A = \text{opt}(c)$ and $B = \text{opt}'(c)$. Since c is canonical, by observation 1, we can assume that $c(B) = 1$. Let c' be a cost function such that $c'(a) = \frac{1}{|A|}$ for $a \in A$, $c'(b) = \frac{1}{|B|}$ for $b \in B$ and $c'(x) = 1$ otherwise. Finally, consider the cost function $c' + \epsilon c$. By lemma 2, the claim follows.

Corollary 3. *For any set system (E, \mathcal{B}) , $\phi'(E, \mathcal{B}) \geq \phi(E, \mathcal{B})$. Moreover, $\phi(E, \mathcal{B}) = 1$ iff $\phi'(E, \mathcal{B}) = 1$.*

Lemma 1 shows that the worst frugality ratio is attained at a canonical cost function, which is very non competitive. Thus it might seem that even when the frugality ratio is large, the VCG mechanism could do well for cost functions we care about. The above corollary however shows that if the frugality ratio is high, there are competitive cost functions where the overpayment is large. On the other hand, a low frugality ratio clearly implies that the overpayment is never large. Thus our definition of frugality is a robust one.

We now note that for any canonical cost function c , the VCG payments have a nice structure. The proof is simple and omitted from this extended abstract.

Proposition 1. *Let c be any canonical cost function. Then*

$$p(e, c) = \min_{T \subseteq \text{opt}'(c) : \text{opt}(c) \setminus \{e\} \cup T \in \mathcal{F}} c(T)$$

3.2 Frugoids

For disjoint sets $A, B \in \mathcal{B}$, and any $Y \subseteq B$, we say that $x \in A$ is *dependent* on Y with respect to A, B if x cannot be replaced in A by some element of $B \setminus Y$, i.e. if for any $y \in B \setminus Y$, $A \setminus \{x\} \cup \{y\}$ is not feasible. Define the *set of dependents* of Y with respect to A, B as $D^{A, B}(Y) = \{x \in A : A \setminus \{x\} \cup \{y\} \text{ is not feasible for any } y \in B \setminus Y\}$.

Call a set system (E, \mathcal{B}) a *frugoid* if for every pair of disjoint sets $A, B \in \mathcal{B}$ and every $Y \subseteq B$, $|D^{A, B}(Y)| \leq |Y|$. The following proposition shows that matroids satisfy the above condition for every pair of bases A and B (not necessarily disjoint).

Proposition 2. *Let (E, \mathcal{B}) be a matroid. Then for any $A, B \in \mathcal{B}$ and any $Y \subseteq B$, $|D^{A, B}(Y)| \leq |Y|$.*

³ We assume for simplicity that there is a cost function attaining the frugality. Similar results would hold if this was not the case.

Proof. From the definition of matroids, it follows that any two base sets have the same cardinality. Also for any base sets S and T and any set $T' \subseteq T$ satisfying $|T'| < |T|$, there exists $S' \subseteq S \setminus T'$ with $|S'| = |S| - |T'|$ such that $T' \cup S'$ is a base set. Now consider any $A, B \in \mathcal{B}$ and let $Y \subseteq B$ be arbitrary. Since $|B \setminus Y| < |B|$ and $A, B \in \mathcal{B}$, there exists $X \subseteq A$ with $|X| = |Y|$ such that $(B \setminus Y) \cup X \in \mathcal{B}$. We shall show that $D^{A,B}(Y) \subseteq X$, from which the claim follows. Consider any $a \in A \setminus X$. Since $|A \setminus \{a\}| < |A|$ and $(B \setminus Y) \cup X \in \mathcal{B}$, there exists a $b \in ((B \setminus Y) \cup X) \setminus (A \setminus \{a\})$ such that $A \setminus \{a\} \cup \{b\}$ is a base set. However, then b must belong to $B \setminus Y$. Thus $a \notin D^{A,B}(Y)$. Since a was arbitrary, $D^{A,B}(Y)$ contains no elements from $A \setminus X$, and hence must be contained in X .

Thus every matroid is a frugoid. We give below an example of a set system which is a frugoid but not a matroid.

Example 1. Consider the set system (E_1, \mathcal{B}_1) where $E_1 = \{a_1, a_2, b_1, b_2, c_1, c_2\}$ and \mathcal{B}_1 contains the following sets:

$$\begin{aligned} & \{a_1, b_1, c_1\} \quad \{a_2, b_1, c_1\} \quad \{a_1, b_2, c_1\} \quad \{a_1, b_1, c_2\} \\ & \{a_2, b_2, c_2\} \quad \{a_1, b_2, c_2\} \quad \{a_2, b_1, c_2\} \quad \{a_2, b_2, c_1\} \\ & \{a_1, a_2\} \end{aligned}$$

Since any pair of disjoint sets in \mathcal{B}_1 come from the matroid $(E_1, \mathcal{B}_1 \setminus \{\{a_1, a_2\}\})$, it is a frugoid.

The following theorem gives an exact characterization of set systems that have frugality ratio at most one.

Theorem 2. *A set system (E, \mathcal{B}) is a frugoid iff $\phi(E, \mathcal{B}) \leq 1$.*

Proof. (proof of \Rightarrow) From Lemma 1, it suffices to show that for all canonical cost functions c , $\phi(E, \mathcal{B}, c) \leq 1$. Let c be a canonical cost function, let $A = \text{opt}(c)$ and $B = \text{opt}'(c)$. We need to show that $p(A) \leq c(B)$. We shall find a one-one mapping π from A to B such that for all $a \in A$, $p(a)$ will be no more than $c(\pi(a))$. The claim would then follow. Consider a bipartite graph with vertex set $A \cup B$, and an edge between a and b if $(A \setminus \{a\}) \cup \{b\}$ is feasible (and hence $p(a) \leq c(b)$). The condition $|D^{A,B}(Y)| \leq |Y|$ implies that there is no Hall set in A (recall that a hall set is a set A such whose neighbourhood is of size strictly smaller than A itself). Hall's theorem [13](see [14] for a proof) then implies that we can find such a mapping. The claim follows.

(Proof of \Leftarrow) We first show that any two disjoint base sets must have the same cardinality. Assume the contrary. Let A and B be disjoint sets with different cardinalities. Without loss of generality, $|A| > |B|$. Consider the cost function c defined as follows:

$$c(e) = \begin{cases} 0 & \text{if } e \in A \\ 1 & \text{if } e \in B \\ \infty & \text{otherwise} \end{cases}$$

Now consider any $e \in A$. Since $A \setminus \{e\}$ is not feasible, $p(e, c) \geq 1$. Thus $p(A, c) \geq |A|$. On the other hand, $c(B) = |B| < |A|$. This contradicts the fact that frugality ratio of (E, \mathcal{B}) is at most 1.

Now suppose the condition is violated for some A, B, Y . Consider the canonical cost function c defined as follows:

$$c(e) = \begin{cases} 0 & \text{if } e \in A \\ 1 & \text{if } e \in B \setminus Y \\ 2 & \text{if } e \in Y \\ \infty & \text{otherwise} \end{cases}$$

Again, for any $e \in A$, $p(e, c) \geq 1$. Moreover, from the definition of dependency, for any $e \in A$ which is 1-dependent on Y , $p(e, c) \geq 2$. Thus $p(A, c) \geq |A| + |D^{A,B}(Y)|$. On the other hand, $c(B) = |B| + |Y|$. Since $|A| = |B|$ and $|D^{A,B}(Y)| > |Y|$, $\phi(E, \mathcal{B}, c) > 1$, which is a contradiction. Hence the claim follows.

3.3 Non frugoids

For systems which are not frugoids, we give some upper bounds on the frugality ratios.

Analogous to the definition of dependence, we define k -dependence as follows. For disjoint sets $A, B \in \mathcal{B}$, and any $Y \subseteq B$, we say that $x \in A$ is k -dependent on Y with respect to A, B if x cannot be replaced in A by at most k elements of $B \setminus Y$, i.e. if for any $X \subseteq B \setminus Y : |X| \leq k, A \setminus \{x\} \cup X$ is not feasible. Define the set of k -dependents of Y with respect to A, B as $D_k^{A,B}(Y) = \{x \in A : A \setminus \{x\} \cup X \text{ is not feasible for any } X \subseteq B \setminus Y, |X| \leq k\}$.

Theorem 3. *The following hold:*

- (i) *Let (E, \mathcal{B}) be a set system such that for every pair of disjoint sets A and B , and for some positive integer k , and for all $Y \subseteq B$, it is the case that $k \frac{|D_k^{A,B}(Y)|}{|Y|} \leq f$. Then the frugality ratio $\phi(E, \mathcal{B}) \leq f$.*
- (ii) *Let (E, \mathcal{B}) be a set system such that the size of each set in \mathcal{B} is at most l . Then $\phi(E, \mathcal{B}) \leq l$.*
- (iii) *Let (E, \mathcal{B}) be a set system instance derived from a set cover problem where each set is of size at most k . Then $\phi(E, \mathcal{B}) \leq k$.*

Proof. The proof of (i) is analogous to theorem 2. (ii) is immediate from the fact that for any $e \in \text{opt}(c)$, $p(e, c) \leq OPT'(c)$. We prove (iii) below. Let c be a canonical cost function. Consider any edge e in $\text{opt}(c)$. Since c is canonical, $p(e, c) \leq c(T_e)$ for any $T_e \subseteq \text{opt}'(c)$ such that $\text{opt}(c) \setminus \{e\} \cup T_e$ is feasible. We construct one such T_e as follows. For each $v' \in e$ that is covered only by e in $\text{opt}(c)$, add to T_e an arbitrary set e' from $\text{opt}'(c)$ that covers v' . We say then that e requires e' to cover v' . We now show that $\sum_{e \in \text{opt}(c)} c(T_e) \leq k OPT'(c)$. Let $e'' = \{v'_1, v'_2, \dots, v'_l\}$ be any set in $\text{opt}'(c)$. If v'_i is covered more than once in $\text{opt}(c)$, no set in $\text{opt}(c)$ requires e'' to cover v'_i . Otherwise, for each v'_i , there is at most one set in $\text{opt}(c)$ that requires e'' to cover v'_i . Hence in all, at most k T_e 's for $e \in \text{opt}(c)$ contain e'' . Now $p(\text{opt}(c), e) \leq \sum_{e \in \text{opt}(c)} c(T_e) = \sum_{e'' \in \text{opt}'(c)} \sum_{e \in \text{opt}(c): e'' \in T_e} c(e'') \leq kc(\text{opt}'(c))$. Hence the claim follows.

The same upper bounds can be shown on the marginal frugality ϕ' as well. We omit the proof from this extended abstract.

We now state some simple lower bounds on the frugality of set systems.

Theorem 4. *The following hold:*

- (i) *Let (E, \mathcal{B}) be a set system such that \mathcal{B} contains two disjoint sets A and B such that $|A| = \alpha|B|$. Then $\phi(E, \mathcal{B}) \geq \alpha$.*
- (ii) *Let (E, \mathcal{B}) be a set system such that \mathcal{B} contains two disjoint sets A and B such that for all $a \in A$, for all $Y \subseteq B$ of cardinality less than k , $A \setminus \{a\} \cup Y$ is not feasible. Then $\phi(E, \mathcal{B}) \geq k \frac{|A|}{|B|}$.*

Proof. Consider the cost function:

$$c(e) = \begin{cases} 0 & \text{if } e \in A \\ 1 & \text{if } e \in B \\ \infty & \text{otherwise} \end{cases}$$

It is easy to check that this cost function gives the lower bounds in both cases.

Note that because of corollary 3, these lower bounds apply to the marginal frugality ϕ' as well.

In the next section, we use these results to estimate the frugality ratio of several important combinatorial problems.

4 Examples

For a class of set system instances, we define the frugality ratio of the class as the largest frugality ratio for an instance from the class. For example, from theorem 2, it follows that the minimum spanning tree problem on graphs (being a matroid) has frugality 1. Figure 1 tabulates some simple consequences of theorems 3 and 4.

Problem	Frugality ratio	Theorems used
Any matroid problem	1	Prop.2 and Thm. 2
Vertex Cover with maximum degree d	d	Thms. 2(ii) and 3(i)
Edge Cover	2	Thms. 2(ii) and 3(ii)
Bipartite graph matching	$\Theta(n)$	Thms. 2(ii) and 3(ii)
Minimum cut	$\Theta(n)$	Thms. 2(ii) and 3(i)
Minimum vertex cut	$\Theta(n)$	Thms. 2(ii) and 3(i)
Dominating set	$\Theta(n)$	Thms. 2(ii) and 3(i)
Set cover - each set of size k	k	Thms. 2(iii) and 3(ii)
Uncapacitated facility location	4	See discussion below

Fig. 1. Frugality ratios of some combinatorial problems

Frugality and locality ratio

We note an interesting connection between frugality and *locality ratio* of a natural local search procedure for the facility location problem. The notion of frugality here is slightly different. We assume that the facilities are owned by agents, and only they know the facility cost. The distances however are well known. Analogous to the locality ratio analysis (theorem 4.3) of Arya et.al. [2], we can show that the payment to the facilities in opt is no more than $(\text{cost}_f(\text{opt}) + 2\text{cost}_f(\text{opt}') + 3\text{cost}_s(\text{opt}'))$ where $\text{cost}_f()$ and $\text{cost}_s()$ denote the facility and the service costs of a solution. Thus the total “expenditure” in this case is $(\text{cost}_f(\text{opt}) + 2\text{cost}_f(\text{opt}') + 3\text{cost}_s(\text{opt}')) + \text{cost}_s(\text{opt}) \leq 4\text{cost}(\text{opt}')$. We omit the details from this extended abstract.

5 Discussion

In this section, we look at some interesting implications of the positive and negative results of the previous sections.

We first look at the problem of computing the VCG payments. In general, this requires solving an optimization problem for finding the optimal solution, and then solving one optimization problem for each agent in the optimal solution. [7] show that whenever the “agents are substitutes” condition holds, all the VCG payments can be computed using the variables in the dual of a linear programming formulation of the underlying optimization problem. (Intuitively, the dual variables correspond to the effect of the corresponding primal constraint on the value of the optimal, which is precisely the “bonus” to the agent in VCG). Since the “agents are substitutes” condition is equivalent⁴ to the set system having frugality ratio 1, the VCG payment for frugoids can be computed by solving a single linear program and its dual.

One criticism of VCG is that it requires all agents to divulge their true values to the auctioneer (and hence trust the auctioneer to not somehow use this information, e.g. in similar cases in the future). A suggested solution to this is to design *iterative mechanisms*, where agents respond to a sequence of offers made by the auctioneer. This has the advantages of simplicity and privacy (see e.g. [1], [23], etc. for further arguments in favour of iterative auctions). [7] also show that under the “agents are substitutes” condition, an iterative mechanism can be designed that gives the same outcome as the VCG mechanism. This, then holds for all frugoids. Moreover [6] conjecture that if the substitutes condition does not hold, there is no iterative mechanism yielding the Vickrey outcome. Assuming this conjecture then, frugoids is exactly the class of minimization problems which have iterative mechanisms implementing the social optimum.

We also note that in general, VCG is the only truthful mechanism that selects the optimal allocation and satisfies individual rationality. Moreover, since every dominant strategy mechanism has an equivalent truthful mechanism (the revelation principle), better frugality ratios cannot be achieved by any mechanism

⁴ We omit the simple proof from this extended abstract

while maintaining optimality. However by imposing some additional restrictions, the class of truthful mechanisms can be made larger and more frugal mechanisms can be designed.

While we have addressed the questions of frugality of exact mechanisms, it turns out that approximate mechanisms need not even be approximately frugal, even for frugoids. In fact, there is a constant factor approximation algorithm for the minimum spanning tree problem that can be implemented truthfully, but the resulting mechanisms has frugality ratio $\Omega(n)$.

6 Conclusion and Further Work

We have defined the frugality ratio which is a robust measure of the economic performance of a mechanism in the presence of competition. We show that frugoids, a natural generalization of matroids is the exact class of problems with frugality ratio 1. We also give lower and upper bounds on the frugality ratio of set systems. We use these to estimate the frugality of several interesting combinatorial problems. An exact characterization of set systems with frugality ratio exactly k for $k > 1$ is an interesting open problem. Further, while it turns out that some important problems have large frugality ratios in the worst case, it would be interesting to see if we can get better positive results by restricting the instances and/or restricting the cost functions in some reasonable way.

We also define the notion of *marginal frugality* and show that it is always lower bounded by the frugality ratio. We leave open the intriguing question of whether or not they are equal.

7 Acknowledgements

I would like to thank Christos Papadimitriou for several useful discussions and comments. I would also like to thank Ranjit Jhala, Amin Saberi and Nikhil Devanur for useful discussions. Many thanks to Amir Ronen, Aaron Archer and Vijay Vazirani for discussions that led to the frugality questions addressed here. I would also like to thank the anonymous referees for their several helpful comments.

References

1. LAWRENCE M. AUSUBEL, "An Efficient Ascending-Bid Auction for Multiple Objects," Working Paper, University of Maryland, Department of Economics, 1997.
2. VIJAY ARYA, NAVEEN GARG, ROHIT KHANDEKAR, ADAM MEYERSON, KAMESH MUNAGALA, VINAYAKA PANDIT, "Local search heuristics for k-median and facility location problems", *Proceedings of the 33rd ACM symposium on the theory of computing, 2001*.
3. AARON ARCHER, CHRISTOS PAPADIMITRIOU, KUNAL TALWAR, ÉVA TARDOS, "An approximate truthful mechanism for combinatorial auctions with single parameter agents", *To appear in SODA 2003*

4. AARON ARCHER, ÉVA TARDOS “Truthful mechanisms for one-parameter agents”, *Proceedings of the 42nd IEEE annual symposium on Foundations of Computer Science*, 2001.
5. AARON ARCHER, ÉVA TARDOS “Frugal Path Mechanisms”, *SODA 2002*.
6. SUSHIL BIKHCHANDANI, SVEN DE VRIES, JAMES SCHUMMER, RAKESH VOHRA “Linear Programming and Vickrey Auctions,” *IMA Volumes in Mathematics and its Applications, Mathematics of the Internet: E-Auction and Markets*, 127:75-116, 2001.
7. SUSHIL BIKHCHANDANI, JOSEPH M. OSTROY, “The package assignment model”, working paper, 2001.
8. E. H. CLARKE “Multipart pricing of public goods”, *Public Choice*, 8:17-33, 1971.
9. EDITH ELKIND, AMIT SAHAI, “Shortest Paths are costly”, *manuscript*, October 2002.
10. JOAN FEIGENBAUM, CHRISTOS PAPADIMITRIOU, AND SCOTT SHENKER, “Sharing the cost of multicast transmissions”, *Journal of Computer and System Sciences* 63:21–41, 2001.
11. THEODORE GROVES “Incentives in teams”, *Econometrica*, 41(4):617-631, 1973.
12. RAHUL GARG, VIJAY KUMAR, ATRI RUDRA AND AKSHAT VERMA, “When can we devise frugal mechanisms for network design problems?”, *manuscript*, 2002.
13. PHILIP HALL, “On representatives of subsets,” *Journal of London Mathematics Society*, 10:26-30, 1935.
14. FRANK HARARY, “Graph Theory”, Addison Wesley, 1971.
15. JOHN HERSHBERGER, SUBHASH SURI, “. Vickrey Pricing in network routing: Fast payment computation”, *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.
16. DANIEL LEHMANN, LIADAN ITA O’CALLAGHAN, YOAV SHOHAM “Truth revelation in rapid approximately efficient combinatorial auctions”, *1st ACM conference on electronic commerce*, 1999
17. AHUVA MU’ALEM, NOAM NISAN “Truthful approximation mechanism for restricted combinatorial auctions”, to appear in *AAAI 2002*
18. ANDREU MAS-COLELL, MICHAEL D. WHINSTON, JERRY R. GREEN“, *Microeconomic Theory*”, Oxford University Press, 1995.
19. NOAM NISAN, AMIR RONEN “Algorithmic mechanism design”, *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing*, 1999.
20. NOAM NISAN, AMIR RONEN “Computationally feasible VCG mechanisms”, *ACM conference on electronic commerce*, 242-252, 2000.
21. MARTIN J. OSBORNE, ARIEL RUBINSTEIN”A course in game theory”, MIT Press, Cambridge, 1994.
22. CHRISTOS PAPADIMITRIOU “Algorithms, games and the Internet”, *Proceedings of the 33rd Annual ACM Symposium on Theory of Computation*, 749-753, 2001.
23. DAVID PARKES, LYLE UNGAR, “Iterative Combinatorial Auctions: Theory and Practice”, *In Proc. 17th National Conference on Artificial Intelligence, (AAAI-00)* pp. 74-81, 2000.
24. WILLIAM VICKREY “Counterspeculation, auctions and competitive sealed tenders”, *Journal of Finance*, 16:8-37, 1961.

A Proof of lemma 1

Lemma 1

Let (E, \mathcal{B}) be a set system, and c be a non negative cost function on E . Then there is a cost function c' such that the following hold:

- c' is canonical.
- $\phi(E, \mathcal{B}, c') \geq \phi(E, \mathcal{B}, c)$

Proof. We change c to a canonical cost function c' . We do so via a sequence of intermediate cost functions, such that the payment to the optimum goes up at each step. Let c'_0 be equal to c on $\text{opt}(c) \cup \text{opt}'(c)$ and ∞ otherwise. Now,

$$\begin{aligned} p(e, c) &= OPT(c[e \mapsto \infty]) - OPT(c[e \mapsto 0]) \\ &= OPT(c[e \mapsto \infty]) - (c(\text{opt}(c)) - c(e)) \\ &\leq OPT(c'_0[e \mapsto \infty]) - (c'_0(\text{opt}(c)) - c'_0(e)) \\ &= p(e, c'_0) \end{aligned}$$

where the equality in step 2 holds because $e \in \text{opt}(c)$ and the inequality in step 3 follows from the fact that $c'_0(e) \geq c(e)$ for all $e \in E$. Now let $\text{opt}(c) = \text{opt}(c'_0) = \{e_1, e_2, \dots, e_k\}$. Define c'_i as follows:

$$c'_i(e) = \begin{cases} 0 & \text{if } e \in \{e_1, \dots, e_i\} \\ c(e) & \text{if } e \in \text{opt}'(c) \cup \{e_{i+1}, \dots, e_k\} \\ \infty & \text{otherwise} \end{cases}$$

It is easy to see that the cost function $c' = c'_k$ is canonical, and that $\text{opt}(c) = \text{opt}(c')$ and $\text{opt}'(c) = \text{opt}'(c')$. Now

$$\begin{aligned} p(e_i, c'_i) &= OPT(c'_i[e_i \mapsto \infty]) - (c'_i(\text{opt}(c)) - c'_i(e_i)) \\ &= OPT(c'_i[e_i \mapsto \infty]) - (c'_{i-1}(\text{opt}(c)) - c'_{i-1}(e_i)) \\ &= OPT(c'_{i-1}[e_i \mapsto \infty]) - (c'_{i-1}(\text{opt}(c'_{i-1})) - c'_{i-1}(e_i)) \\ &= OPT(c'_{i-1}[e_i \mapsto \infty]) - OPT(c'_{i-1}[e_i \mapsto 0]) \\ &= p(e_i, c'_{i-1}) \end{aligned}$$

and

$$\begin{aligned} p(e_j, c'_i) &= OPT(c'_i[e_j \mapsto \infty]) - OPT(c'_i[e_j \mapsto 0]) \\ &= OPT(c'_i[e_j \mapsto \infty]) - (c'_i(\text{opt}(c)) - c'_i(e_j)) \\ &= OPT(c'_i[e_j \mapsto \infty]) - ((c'_{i-1}(\text{opt}(c)) - c(e_i)) - c'_i(e_j)) \\ &= OPT(c'_i[e_j \mapsto \infty]) + c(e_i) - (c'_{i-1}(\text{opt}(c'_{i-1})) - c'_{i-1}(e_j)) \\ &\geq OPT(c'_{i-1}[e_j \mapsto \infty]) - (c'_{i-1}(\text{opt}(c'_{i-1})) - c'_{i-1}(e_j)) \\ &= p(e_j, c'_{i-1}) \end{aligned}$$

where the inequality follows from the fact that $\text{opt}(c'_i[e_j \mapsto \infty])$ is a feasible set with cost $OPT(c'_i[e_j \mapsto \infty]) + c(e_i)$ under c'_{i-1} .

Thus $p(\text{opt}(c'_i), c'_i) \geq p(\text{opt}(c'_{i-1}), c'_{i-1})$. Thus $p(\text{opt}(c'), c') \leq p(\text{opt}(c), c)$. Also $OPT'(c) = OPT'(c')$. Hence the claim follows.