False-Name-Proof Mechanisms for Hiring a Team

Atsushi Iwasaki[1], David Kempe[2], Yasumasa Saito[1],
Mahyar Salek[2], and Makoto Yokoo[1]

[1] Department of ISEE, Kyushu University, Fukuoka 819-0395, Japan,
{iwasaki@, saito@agent, yokoo@}is.kyushu-u.ac.jp
[2] Department of Computer Science, University of Southern California, CA
90089-0781, USA, {dkempe, salek}@usc.edu

Abstract. We study the problem of hiring a team of selfish agents to perform a task. Each agent is assumed to own one or more elements of a set system, and the auctioneer is trying to purchase a feasible solution by conducting an auction. Our goal is to design auctions that are truthful and false-name-proof, meaning that it is in the agents' best interest to reveal ownership of all elements (which may not be known to the auctioneer a priori) as well as their true incurred costs.We first propose and analyze a false-name-proof mechanism for the special cases where each agent owns only one element in reality. We prove that its frugality ratio is bounded by $n2^n$, which nearly matches a lower bound of $\Omega(2^n)$ for all false-name-proof mechanisms in this scenario. We then propose a second mechanism. It requires the auctioneer to choose a reserve cost a priori, and thus does not always purchase a solution. In return, it is false-name-proof even when agents own multiple elements. We experimentally evaluate the payment (as well as social surplus) of the second mechanism through simulation.

1 Introduction

In the problem of *hiring a team of agents* [1–3], an auctioneer knows which subsets of agents can perform a complex task together, and needs to hire such a team (called a *feasible set* of agents). Since the auctioneer does not know the true costs incurred by agents, we assume that the auctioneer will use an auction to elicit bids. A particularly well-studied special case of this problem is that of a *path auction* [1, 4–6]: the agents own edges of a known graph, and the auctioneer wants to purchase an *s*-*t* path.

Selfish agents will try to maximize their profit, even if it requires misrepresenting their incurred cost or their identity. Thus, the auctioneer should design the auction to be *truthful*, i.e., making it in agents' best interest to reveal actual costs and ownership. The area of designing such auctions is known as *mechanism design* [6–8]. Most recent results on truthful mechanism design have focused on discouraging misrepresentation of costs. However, as recently pointed out by Yokoo et al. in the context of combinatorial auctions [9, 10], a second threat is that of *false-name manipulations*, in which agents owning multiple elements of the underlying set system invent "pseudo-agents" in order to pretend that all these agents must be paid, leading to higher total payments.

1.1 Our contributions

We introduce a model of false-name manipulation in auctions for hiring a team, such as s-t path auctions. In this model, the set system structure and element ownership are not completely known to the auctioneer. Thus, in order to increase profit, an agent who owns an element can pretend that the element is in fact a set consisting of multiple elements owned by different agents. Similarly, an agent owning multiple elements can submit bids for these elements under different identities. We call a mechanism *false-name-proof* if it is truthful, and a dominant strategy is for each agent to reveal ownership of all elements.

Our first main contribution is a false-name-proof mechanism MP for the special case in which each agent owns exactly one element. This mechanism introduces an exponential multiplicative penalty against sets in the number of participating agents. We show that its frugality ratio (according to the definition of Karlin et al. [5]) is at most $n2^n$ for all set systems of n elements, which nearly matches a worst-case lower bound of $\Omega(2^n)$ we establish for *every* false-name-proof mechanism.

When agents may own multiple elements, we present an alternative mechanism AP, based on an a priori chosen reserve cost r and additive penalties. The mechanism is false-name-proof in the general setting, but depends crucially on the choice of r, as it will not purchase a solution unless there is one whose cost (including the penalty) is at most r. We investigate the AP mechanism experimentally for s-t path auctions on random graphs, observing that AP provides social surplus not too far from a Pareto-efficient one at an appropriate reserve cost.

1.2 Related Work

If false-name bids are not a concern, then it has long been known that the VCG mechanism gives a truthful mechanism and identifies the Pareto optimal solution. As the payments of VCG can be significantly higher than the cheapest alternative solution, several papers [1, 3–5] have investigated the *frugality* of mechanisms: the overpayment compared to a natural lower bound. In particular, [5] presents a mechanism called the $\sqrt{}$ mechanism achieving frugality ratio within a constant factor of optimal for s-t path auctions in graphs.

The issue of false-name bids was recently studied in combinatorial auctions and several special cases by Yokoo et al. [11–14, 9], who developed false-name-proof mechanisms in those scenarios, but also proved that no mechanism can be both false-name-proof and Pareto efficient. Notice that the false-name-proof mechanisms for combinatorial procurement auctions given in [12, 13] cannot be applied in our setting, as they assume additive valuations on the part of the auctioneer, i.e., that the auctioneer derives partial utility from partial solutions.

A somewhat similar scenario arises in job scheduling, where users may split or merge jobs to obtain earlier assignments. Moulin [15] gives a mechanism that is strategy-proof against both merges and splits and achieves efficiency within a constant factor of optimum. However, when agents can exchange money, no such mechanism is possible [15].

For the specific case of path auctions, the impact of false-name bids was recently studied by Du et al. [16]. They showed that if agents can own multiple edges, then there is no false-name-proof and efficient mechanism. Furthermore, if bids are anonymous, i.e., agents do not report any identity for edge ownership, then no mechanism can be strategy-proof. Notice that this does not preclude false-name-proof and truthful mechanisms in which the auctioneer takes ownership of multiple edge by the same agent into account, and rewards the agent accordingly.

2 Preliminaries

Our framework is based on that of [1, 17, 5, 3]. A *set system* $(E, \mathcal{F})$ is specified by a set E of n *elements* and a collection $\mathcal{F} \subseteq 2^E$ of *feasible sets*. For instance, in the important special case of an s-t path auction, $S \in \mathcal{F}$ if and only if S is an s-t path.

Agents can own multiple elements, and A^i denotes an element of a partition $\mathcal{A}$ of E and the set of elements owned by agent i. An *owned set system*, i.e., a set system with ownership structure, is specified by $((E, \mathcal{F}), \mathcal{A})$. Each element e has an associated *cost* c_e, the true cost that its owner $o(e)$ will incur if e is selected by the mechanism.[3] This cost is *private*, i.e., known only to $o(e)$. An *auction* consists of two steps:

1. Each agent i submits sealed bids $(b_e, \tilde{o}(e))$ for elements e, where $\tilde{o}(e)$ denotes the identifier of e's purported owner (which need not be the actual owner).
2. Based on the bids, the auctioneer uses an algorithm that is common knowledge among the agents in order to select a feasible set $S^* \in \mathcal{F}$ as the winner and compute a payment p_i for each agent i with an element e such that $i = \tilde{o}(e)$. We say that the elements $e \in S^*$ *win*, and all other elements *lose*.

The *profit* of an agent i is the sum of all payments she receives, minus the incurred cost $c(S^* \cap A^i)$. Each agent is only interested in maximizing her profit, and might choose to misrepresent ownership or costs to this end. However, we assume that agents do not collude. If agents report correct ownership for all $e \in A^i$, then a mechanism is truthful by definition if for any fixed vector b^{-i} of bids by all agents other than i, it is in agents i's best interest to bid $b_e = c_e$ for all $e \in A^i$, i.e., agent e's profit is maximized by bidding $b_e = c_e$ for all these elements e.

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

In this paper, we extend the study of truthful mechanisms to take into account *false-name manipulation*: agents claiming ownership of non-existent elements (which we call *self-division*) or choosing not to disclose ownership of elements (which we call *identifier splitting*).

Definition 1 (Identifier Splitting [9, 10]). *An agent i owning a set A^i may choose to use different identifiers in her bid for some or all of the elements. Formally, the owned set system $((E, \mathcal{F}), \mathcal{A})$ is replaced by $((E, \mathcal{F}), \mathcal{A}')$, where $A^{'} = A \setminus \{A^i\} \cup \{A^{i'}\} \cup \{A^{i''}\}$, and $A^i = A^{i'} \cup A^{i''}$.*

Definition 2 (Self-Division). *An agent i owning element e is said to* self-divide *e if e is replaced by two or more elements $e_1, \ldots, e_k$, and different owners are reported for the e_i. Formally, the owned set system $((E, \mathcal{F}), \mathcal{A})$ is replaced by $((E', \mathcal{F}'), \mathcal{A}')$, whose elements are $E' = E \setminus \{e\} \cup \{e_1, \ldots, e_k\}$, such that the feasible sets $\mathcal{F}'$ are exactly those sets S not containing e, as well as sets $S \setminus \{e\} \cup \{e_1, \ldots, e_k\}$ for all feasible sets $S \in \mathcal{F}$ containing e. The ownership structure is $A^{i_j} = \{e_j\}$ for $j = 1, \ldots, k$, where each i_j is a new agent.*

Intuitively, self-division allows an agent to pretend that multiple distinct agents are involved in doing the work of element e, and that each of them must be paid separately. For self-division to be a threat, there must be uncertainty on the part of the auctioneer about the true set system $(E, \mathcal{F})$. In particular, it is meaningless to talk about a mechanism for an individual set system, as the auctioneer does not know a priori what the set system is. Hence, we define *classes of set systems closed under subdivision*, as the candidate classes on which mechanisms must operate.

Definition 3.
1. *For two set systems $(E, \mathcal{F})$ and $(E', \mathcal{F}')$, we say $(E', \mathcal{F}')$ is* reachable *from $(E, \mathcal{F})$ by subdivisions if $(E', \mathcal{F}')$ is obtained by (repeatedly) replacing individual elements $e \in E$ with $\{e_1, \ldots, e_k\}$, such that the feasible sets $\mathcal{F}'$ are exactly those sets S not containing e, as well as sets $S \setminus \{e\} \cup \{e_1, \ldots, e_k\}$ for all feasible sets $S \in \mathcal{F}$ containing e.*
2. *A class $\mathcal{C}$ of set systems is* closed under subdivisions *iff with $(E, \mathcal{F})$, all set systems reachable from $(E, \mathcal{F})$ by subdivisions are also in $\mathcal{C}$.*

For example, *s*-*t* path auction set systems are closed under subdivisions, whereas minimum spanning tree set systems are not (because subdivisions would introduce new nodes that must be spanned).

In both identifier splitting and self-division, we will sometimes refer to the new agents i' whose existence i invents as *pseudo-agents*. A mechanism is *false-name-proof* if it is a dominant strategy for each agent i to simply report the pair (c_e, i) as a bid for each element $e \in A^i$. Thus, neither identifier splitting nor self-division nor bids $b_e \neq c_e$ can increase the agent's profit. Among other things, this allows us to use b_e and c_e interchangeably when discussing false-name-proof mechanisms. Notice that we explicitly define the concept of false-name-proof mechanisms to imply that the mechanism is also truthful when each agent i owns only one element.

Efficiency and Frugality

A mechanism is *Pareto efficient* if it always maximizes the sum of all participants' utilities (including that of the auctioneer). While it is well known that the VCG mechanism is truthful and Pareto efficient, Du et al. [16] show that there is no Pareto efficient and false-name-proof mechanism, even for *s*-*t* path auctions. Yokoo et al. [10] showed the same for combinatorial auctions.

While Pareto efficient mechanisms maximize social welfare, they can significantly overpay compared to other mechanisms [5]. In order to analyze the overpayment, we use the definition of *frugality ratio* from [5].

Definition 4 ([5]). *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). For any vector of costs* $\mathbf{c}$ *for elements, 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) $b(S \setminus T) \leq c(T \setminus S)$ *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* $b(S \setminus T_e) = c(T_e \setminus S)$

This definition essentially captures the payments in a "cheapest Nash Equilibrium" of a first-price auction, and gives a natural lower bound generalizing second-lowest cost for comparison purposes.

Definition 5. *The frugality of a mechanism* $\mathcal{M}$ *for a set system* $(E, \mathcal{F})$ *is*

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

i.e., the worst case, over all cost vectors $\mathbf{c}$*, of the overpayment compared to the "first-price" payments. Here,* $p_{\mathcal{M}}(\mathbf{c})$ *denotes the total payments made by* $\mathcal{M}$ *when the cost vector is* $\mathbf{c}$*.*

3 A Multiplicative Penalty Mechanism

We present a mechanism MP based on exponential multiplicative penalties. It is false-name-proof for arbitrary classes of set systems closed under subdivisions, *so long as each agent only owns one element*. We can therefore identify elements e with agents. Since we assume each agent owns exactly one element, $\mathcal{A}$ is automatically determined by E, so we can focus on set systems instead of owned set systems. After the agents submit bids b_e for elements, MP chooses the set S^* minimizing $b(S) \cdot 2^{|S|-1}$, among all feasible sets $S \in \mathcal{F}$. Each agent $e \in S^*$ is then paid her threshold bid $2^{|S^{-e}|-|S^*|} b(S^{-e}) - b(S^* \setminus \{e\})$, where S^{-e} denote the best solution (with respect to the objective function $b(S) \cdot 2^{|S|-1}$) among feasible sets S not containing e. Notice that while this selection may be NP-hard in general, it can be accomplished in polynomial time for path auctions, by using the Bellman/Ford algorithm to compute the shortest path for each number of hops, and then comparing among the at most n such shortest paths.

Theorem 1. *For all classes of set systems closed under subdivision,* MP *is false-name-proof, so long as each agent only owns one element. Furthermore, it has frugality ratio $O(n \cdot 2^n)$, where $n = |E|$.*

Proof. If an agent $e = e_0$ self-divides into $k+1$ elements $e_0, \ldots, e_k$, then either all of the e_i or none of them are included in any feasible set S. Thus, we can always think of just one threshold $\tau_k(e)$ for the self-divided agent e: if the sum of the bids of all the new elements e_j exceeds $\tau_k(e)$, then e loses; otherwise, it is paid at most $(k+1)\tau_k(e)$. The original threshold of agent e is $\tau(e) = \tau_0(e)$.

The definition of the MP mechanism implies $\tau_k(e) \leq 2^{-k}\tau(e)$. If e still wins after self-division (otherwise, there clearly is no incentive to self-divide), the total payment to e is at most $(k+1)2^{-k}\tau(e)$. The alternative of not self-dividing, and submitting a bid of 0, yields a payment of $\tau(e) \geq (k+1)2^{-k}\tau(e)$. Thus, refraining from self-division is a dominant strategy. Given that no agent will submit false-name bids, the monotonicity of the selection rule implies that the mechanism is incentive compatible, and we can assume that $b_e = c_e$ for all agents e.

To prove the upper bound on the frugality ratio, consider again any winning agent $e \in S^*$. Her threshold bid is $\tau(e) = \min_{T \in \mathcal{F}: e \notin T} 2^{|T|-|S^*|}c(T) - c(S^* \setminus \{e\})$, and the total payment is the sum of individual thresholds for S^*,

$$\begin{aligned} p_{\text{MP}}(\mathbf{c}) &= \textstyle\sum_{e \in S^*} \min_{T \in \mathcal{F}: e \notin T} 2^{|T|-|S^*|}c(T) - c(S^* \setminus \{e\}) \\ &\leq 2^n \textstyle\sum_{e \in S^*} \min_{T \in \mathcal{F}: e \notin T} c(T). \end{aligned}$$

Let S be the cheapest solution with respect to the c_e, i.e., without regard to the sizes of the sets. By Definition 4, $\nu(\mathbf{c}) = \sum_{e \in S} b_e$, subject to the constraints of the mathematical program given. Focusing on any fixed agent e', we let $T_{e'}$ denote the set from the third constraint of Definition 4, and can rewrite

$$\nu(\mathbf{c}) = \textstyle\sum_{e \in S - T_{e'}} b_e + \sum_{e \in S \cap T_{e'}} b_e = \sum_{e \in T_{e'} - S} c_e + \sum_{e \in T_{e'} \cap S} b_e \geq c(T_{e'}).$$

Since this inequality holds for all e', we have proved that $\nu(\mathbf{c}) \geq \max_{e \in S} c(T_e)$. On the other hand we can further bound the payments by

$$\begin{aligned} 2^n \textstyle\sum_{e \in S^*} \min_{T \in \mathcal{F}: e \notin T} c(T) &\leq n2^n \max_{e \in S^*} \min_{T \in \mathcal{F}: e \notin T} c(T) \\ &\leq n2^n \max_{e \in S} \min_{T \in \mathcal{F}: e \notin T} c(T) \\ &\leq n2^n \max_{e \in S} c(T_e). \end{aligned}$$

Here, the second-to-last inequality followed because for all $e \in S^* \setminus S$, the minimizing set T is actually equal to S, and therefore cannot have larger cost than $c(T_e)$ for any $e \in S$, by definition of S. Thus, the frugality ratio of MP is

$$\phi_{\text{MP}} = \sup_{\mathbf{c}} \tfrac{p_{\text{MP}}(\mathbf{c})}{\nu(\mathbf{c})} \leq \tfrac{n2^n \max_{e \in S} c(T_e)}{\max_{e \in S} c(T_e)} = n2^n. \quad \blacksquare$$

3.1 An Exponential Lower Bound

An exponentially large frugality ratio is not desirable. Unfortunately, any mechanism which is false-name-proof will have to incur such a penalty, as shown by the following theorem.

Theorem 2. *Let $\mathcal{C}$ be any class of monopoly free set systems closed under subdivisions, and $\mathcal{M}$ be any truthful and false-name-proof mechanism for $\mathcal{C}$. Then, the frugality ratio of $\mathcal{M}$ on $\mathcal{C}$ is $\Omega(2^n)$ for set systems with $|E| = n$.*

Proof. Let $(E_0, F_0) \in \mathcal{C}$ be a set system minimizing $|E_0|$. Let $S^* \in F_0$ be the winning set under $\mathcal{M}$ winning when all agents $e \in E_0$ bid 0, and let $e \in S^*$ be arbitrary, but fixed. Because (E_0, F_0) is monopoly free, there must be a feasible set $T \in F_0$ with $e \notin T$ and $T \not\subseteq S^*$. Among all such sets T, let T_e be the one minimizing $|S^* \cup T|$, and let $\hat{e}$ in T_e be arbitrary. Define $Z = T_e \cup S^* \setminus \{e, \hat{e}\}$ (the "zero bidders"), and $I = E_0 \setminus (T_e \cup S^*)$ (the "infinity bidders"). Consider the following bid vector: both e and $\hat{e}$ bid 1, all agents $e' \in Z$ bid 0, and all agents $e' \in I$ bid ∞ . Let W be the winning set. We claim that W must contain at least one of e and $\hat{e}$ (w.l.o.g., assume that $e \in W$). For W cannot contain any of the infinity bidders. And if it contained neither e nor $\hat{e}$, then W would have been a candidate for T_e with smaller $|W \cup S^*|$, which would contradict the choice of T_e.

Now, let (E_k, F_k) be the set system resulting if agent e self-divides into new agents $e_0, \ldots, e_k$, for $k \geq 0$. Define $\tau(j, k)$, for $j = 0, \ldots, k$, to be the threshold bid under $\mathcal{M}$ for agent e_j in the set system (E_k, F_k), given that all $e' \in Z$ bid 0, all $e' \in I$ bid ∞, and all e_i for $i \neq j$ also bid 0, while $\hat{e}$ bids 1. Above, we thus showed that $1 \leq \tau(0, 0) < \infty$. We now show by induction on d that for all d, there exists an $h \leq d$ such that

$$2^{-d} \sum_{i=0}^{k} \tau(i, k) \geq \sum_{i=h}^{k+h} \tau(i, k+d).$$

The base case $d = 0$ is trivial. For the inductive step, assume that we have proved the statement for d. Because $\mathcal{M}$ is truthful, the payment of an agent is exactly equal to the threshold bid, so each agent i is paid $\tau(i, k+d)$ in the auction on the set system (E_{k+d}, F_{k+d}) with the bids as given above. If agent i were to self-divide into two new agents, the new set system would be (E_{k+d+1}, F_{k+d+1}), and the payment of agent i (who is now getting paid as two pseudo-agents i and $i+1$) would be $\tau(i, k+d+1) + \tau(i+1, k+d+1)$. Because $\mathcal{M}$ was assumed to be false-name-proof, it is not in the agent's best interest to self-divide in such a way, i.e., $\tau(i, k+d) \geq \tau(i, k+d+1) + \tau(i+1, k+d+1)$. Summing this inequality over all agents $i = h, \ldots, h+k$, we obtain

$$\begin{aligned} \sum_{i=h}^{h+k} \tau(i, k+d) &\geq \sum_{i=h}^{h+k} (\tau(i, k+d+1) + \tau(i+1, k+d+1)) \\ &= \sum_{i=h}^{h+k} \tau(i, k+d+1) + \sum_{i=h+1}^{h+k+1} \tau(i, k+d+1). \end{aligned}$$

Define $\ell = 0$ if $\sum_{i=h}^{h+k} \tau(i, k+d+1) \leq \sum_{i=h+1}^{h+k+1} \tau(i, k+d+1)$; otherwise, let $\ell = 1$. Then, the above inequality implies that

$$\sum_{i=h}^{h+k} \tau(i, k+d) \geq 2 \sum_{i=h+\ell}^{h+k+\ell} \tau(i, k+d+1).$$

Finally, setting $h' := h + \ell$, we can combine this inequality with the induction hypothesis to obtain that

$$2^{-(d+1)} \sum_{i=0}^{k} \tau(i, k) \geq \sum_{i=h'}^{k+h'} \tau(i, k+d+1),$$

which completes the inductive proof.

Applying this equation with $k = 0$, we obtain that for each $d \geq 0$, there exists an $h \leq d$ such that $\tau(h,d) \leq 2^{-d} \cdot \tau(0,0)$. Thus, in the set system (E_d, F_d), if all infinity bidders have cost ∞, agent h has cost just above $2^{-d}\tau(0,0)$, and all other agents have cost 0, then agent $\hat{e}$ must be in the winning set, and must be paid at least 1. But it is easy to see that in this case, $\nu(c) = 2^{-d}\tau(0,0)$, and the frugality ratio is thus at least $2^d/\tau(0,0) = \Omega(2^d)$ (since $\tau(0,0)$ is a constant independent of d). Finally, $|E_d| = |Z| + |I| + d + 1$, and because Z and I are constant for our class of examples, the frugality ratio is $2^{-(|Z|+|I|-1)} \cdot 2^n/\tau(0,0) = \Omega(2^n)$. ∎

4 An Additive Penalty Mechanism with Reserve Cost

We next propose and analyze a mechanism called AP, based on additive penalties and a reserve cost. It will only purchase a solution when the total cost (including penalties) does not exceed the a priori chosen reserve cost r, and thus requires a judicious choice of r by the auctioneer. In return, AP is false-name-proof even when agents own multiple elements.

For any set $S \in \mathcal{F}$, let $w(S)$ denote the number of (pseudo-)agents owning one or more elements of S, called the *width* of the set S. The width-based *penalty* for a set S of width w is $D_r(w) = \frac{2^{w-1}-1}{2^{w-1}} \cdot r$. Based on the actual costs and the penalty, we define the *adjusted cost* of a set S to be $\beta(S) = b(S) + D_r(w(S))$.

The AP mechanism first determines the set S^* minimizing the adjusted cost $\beta(S)$, among all feasible sets $S \in \mathcal{F}$. If its adjusted cost exceeds the reserve cost r, then AP does not purchase any set, and does not pay any agents. Otherwise, it chooses S^*, and pays each winning agent (i.e., each agent i with $S^* \cap A^i \neq \emptyset$) her threshold bid $p_i = \min(r, \beta(S^{-i})) - (b(S^* \setminus A^i) + D_r(w(S^*)))$ with respect to $\beta(S)$. Here, S^{-i} denotes the best solution with respect to $\beta(S)$ such that S^{-i} contains no elements from A^i.

4.1 Analysis of AP

In this section, we prove that simply submitting the pair (b_e, i) for each element $e \in A^i$ is a dominant strategy for each agent i under the mechanism AP. Furthermore, we prove that the payments of the AP mechanism never exceed r. As a first step, we prove that it never increases an agent's profit to engage in identifier splitting.

Lemma 1. *Suppose that agent i owns elements A^i, and splits identifiers into i', i'', with sets $A^{i'}, A^{i''}$, such that $A^{i'} \cup A^{i''} = A^i$. Then, the profit agent i obtains after splitting is no larger than that obtained before splitting.*

Proof. Let $S^* \in \mathcal{F}$ be the winning set prior to agent i's identifier split. We first consider the case when the winning set does not change due to the identifier split. If only one of the new pseudo-agents i', i'' wins (say, i'), then $\beta(S^{-i'}) \leq \beta(S^{-i})$, because every feasible set not using elements from A^i also does not use elements from $A^{i'}$. Hence, the payment of i could only decrease, and we may henceforth assume that both i' and i'' win, which means that the width of the winning set S^* increases from w to $w+1$.

For simplicity, we write $B^{-i} = \min(r, \beta(S^{-i}))$, and similarly for i' and i''. The payment to i before the split is $B^{-i} - (b(S^* \setminus A^i) + D_r(w))$, whereas the new payment after the split is

$$\begin{aligned}&B^{-i'} - (b(S^* \setminus A^{i'}) + D_r(w+1)) + B^{-i''} - (b(S^* \setminus A^{i''}) + D_r(w+1))\\&= B^{-i'} + B^{-i''} - 2b(S^*) + b(S^* \cap A^i) - 2D_r(w+1).\end{aligned}$$

As argued above, we have that $B^{-i''} \leq B^{-i}$, and by definition of $B^{-i'}$, we also know that $B^{-i'} \leq r$. Thus, canceling out penalty terms, the increase in payment to agent i is bounded from above by

$$B^{-i'} + B^{-i''} - B^{-i} - b(S^*) - r \leq r + B^{-i} - B^{-i} - b(S^*) - r = -b(S^*) \leq 0.$$

Hence, identifier splitting can only lower the payment of agent i. Since the total cost incurred by agent i stays the same, this proves that there is no benefit in identifier splitting.

Next, suppose that the winning set after the split changes to $S'^* \neq S^*$. Clearly, if i does not win at all after the split, i.e., $S'^* \cap A^i = \emptyset$, then i has no incentive to split identifiers. Otherwise, if i does win after the split, then i must also win before the split. For the split can only increase $D_r(w(S))$ for all sets S containing any of i's elements, while not affecting $D_r(w(S))$ for other sets. We can assume w.l.o.g. that agent i bids ∞ on all elements $e \in A^i \setminus S'^*$. For the winning set will stay the same, because $\beta(S'^*)$ stays the same, and $\beta(S)$ can only increase for other sets S, and the payments can only increase.

But then, S'^* will also be the winning set if i does not split identifiers (the adjusted cost $\beta(S'^*)$ decreases, while all other adjusted costs stay the same). Now, we can apply the argument from above to show that the payments to agent i do not increase as a result of splitting identifiers. Thus, so long as an agent can submit bids of false cost instead, it is never a dominant strategy to split identifiers. ■

Theorem 3. *For all classes of set systems closed under subdivision,* AP *is false-name-proof, even if agents can own multiple elements and split identifiers. Thus, for each agent i, submitting bids (c_e, i) for each element $e \in A^i$ is a dominant strategy.*

Proof. First, notice that if an agent owns two elements in the winning solution, AP does not treat the agent differently from if she only owned one element. Thus, the proof of Lemma 1 also shows that self-division can never be beneficial for an agent, and we can assume from now on that no agent will self-divide or split identifiers. Thus, each agent i submits bids (b_e, i) for all elements $e \in A^i$. If the set $S^* \in \mathcal{F}$ wins under AP, agent i's utility is

$$p_i - c(S^* \cap A^i) = B^{-i} - (b(S^* \setminus A^i) + D_r(w(S^*)) + c(S^* \cap A^i)).$$

Since B^{-i} is a constant independent of the bids $b(e)$ by agent i, agent i's utility is maximized when $(b(S^* \setminus A^i) + D_r(w(S^*)) + c(S^* \cap A^i))$ is minimized. But this is exactly the quantity that AP will minimize when agent i submits truthful bids for all her elements; hence, truthfulness is a dominant strategy. ■

The next theorem proves that an auctioneer with a reserve cost of r faces no loss.

Theorem 4. *The sum of the payments made by* AP *to agents never exceeds* r.

Proof. Because we already proved that AP is false-name-proof, we can without loss of generality identify $c(e)$ and $b(e)$ for each element e. When w agents are part of the winning set S^*, the payment to agent i is

$$p_i = B^{-i} - (c(S^* \setminus A^i) + D_r(w)) \leq r - (c(S^* \setminus A^i) + r - \tfrac{r}{2^{w-1}}) = \tfrac{r}{2^{w-1}}$$

Thus, the sum of all payments to agents i is at most $w \cdot \frac{r}{2^{w-1}} \leq r$. ■

4.2 Experiments

Since the AP mechanism does not always purchase a feasible set, we cannot analyze its frugality ratio in the sense of Definition 5. (The definition is based on the assumption that the mechanism always purchases a set.) Instead, we complement the analysis of the previous section with experiments for shortest s-t path auctions on random graphs. Our simulation compares the payments of the AP mechanism with VCG, under the assumption that there is in fact no false-name manipulation and each agent owns one edge. Thus, we evaluate the overpayment caused by preventing false-name manipulation.

Since some of our graphs have monopolies, we modify VCG by introducing a reserve cost r. Thus, if S^* is the cheapest solution with respect to the cost, the reserve-cost VCG mechanism (RVCG) only purchases a path when $c(S^*) \leq r$. In that case, the payment to each edge $e \in S^*$ is $p_e = \min(r, c(S^{-e})) - c(S^* \setminus \{e\})$, where S^{-e} is the cheapest solution not containing e.

Our generation process for random graphs is as follows: 40 nodes are placed independently and uniformly at random in the unit square $[0, 1]^2$. Then, 200 independent and uniformly random node pairs are connected with edges.[4] The cost of each edge e is its Euclidean length. We evaluate 100 random trials; in each, we seek to buy a path between two randomly chosen nodes. While the number of nodes is rather small compared to the real-world networks on which one would like to run auctions, it is dictated by the computational complexity of the mechanisms we study. Larger-scale experiments are a fruitful direction for future work.

Figure 1 shows the average social surplus (the difference between the reserve cost and the true cost incurred by edges on the chosen path, $r - \sum_{e \in S^*} c_e$) in AP

[4] We also ran simulations on random small-world networks [18]. Our results for small-world networks are qualitatively similar, and we therefore focus on the case of uniformly random networks here.

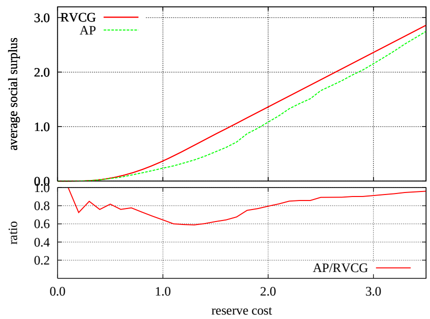

Fig. 1. The evaluation results of social surplus.

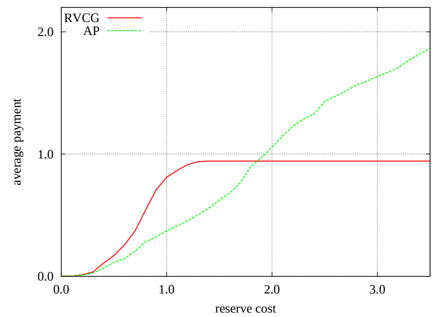

Fig. 2. The evaluation results of payments.

and RVCG, as well as the ratio between the two, when varying the reserve cost $r \in [0, 3.5]$. The social surplus for both increases roughly linearly under both mechanisms. While the plot shows some efficiency loss by using AP, it is always within a factor of about 60% for our instances, and on average around 80%.

Figure 2 illustrates the average payments of the auctioneer. Clearly, small reserve costs lead to small payments, and when the reserve costs are less than 1.8, the payment of AP is in fact smaller than that of RVCG. As the reserve cost r increases, RVCG's payments converge, while those of AP keep increasing almost linearly. The reason is that the winning path in AP tends to have fewer edges than other competing paths, and is thus paid an increased bonus as r increases. We would expect such behavior to subside as there are more competing paths with the same number of edges.

5 Open Questions

It remains open whether there is a mechanism which always purchases a solution, and is false-name-proof even when each agent has multiple elements. This holds even for such seemingly simple cases as s-t path auctions. It may be possible that no such mechanism exists, which would be an interesting result in its own right. The difficulty of designing false-name-proof mechanisms for hiring a team is mainly due to a lack of useful characterization results for incentive-compatible mechanisms when agents have multiple parameters. While a characterization of truthful mechanisms has been given by Rochet [19], this condition is difficult to apply in practice.

It would also be desirable to get the bounds in Section 3 to match asymptotically, i.e., to either remove the factor n from the upper bound, or tighten the lower bound accordingly. The latter may be difficult, as it is likely at least as difficult as designing a truthful mechanism for all set systems with frugality ratio within a constant factor of optimum. Thus, even progress on this question for specific classes of set systems would be desirable.

References

1. Archer, A., Tardos, E.: Frugal path mechanisms. In: Proc. 13th ACM Symp. on Discrete Algorithms, ACM/SIAM (2002) 991–999
2. Garg, R., Kumar, V., Rudra, A., Verma, A.: Coalitional games on graphs: core structures, substitutes and frugality. Technical Report TR-02-60, UTCS (2002)
3. Talwar, K.: The price of truth: Frugality in truthful mechanisms. In: Proc. 21st Annual Symp. on Theoretical Aspects of Computer Science. (2003)
4. Elkind, E., Sahai, A., Steiglitz, K.: Frugality in path auctions. In: Proc. 15th ACM Symp. on Discrete Algorithms, ACM/SIAM (2004)
5. Karlin, A., Kempe, D., Tamir, T.: Beyond VCG: Frugality of truthful mechanisms. In: Proc. 46th IEEE Symp. on Foundations of Computer Science. (2005)
6. Nisan, N., Ronen, A.: Algorithmic mechanism design. In: Proc. 31st ACM Symp. on Theory of Computing, ACM Press (1999) 129–140
7. Mas-Collel, A., Whinston, W., Green, J.: Microeconomic Theory. Oxford University Press (1995)
8. Papadimitriou, C.: Algorithms, games and the internet. In: Proc. 33rd ACM Symp. on Theory of Computing, ACM Press (2001) 749–752
9. Yokoo, M., Sakurai, Y., Matsubara, S.: Robust Combinatorial Auction Protocol against False-name Bids. Artificial Intelligence **130**(2) (2001) 167–181
10. Yokoo, M., Sakurai, Y., Matsubara, S.: The effect of false-name bids in combinatorial auctions: New fraud in Internet auctions. Games and Economic Behavior **46**(1) (2004) 174–188
11. Iwasaki, A., Yokoo, M., Terada, K.: A Robust Open Ascending-price Multi-unit Auction Protocol against False-name bids. Decision Support Systems **39**(1) (March 2005) 23–39
12. Suyama, T., Yokoo, M.: Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction. Autonomous Agents and Multi-Agent Systems **11**(1) (2005) 7–21
13. Suyama, T., Yokoo, M.: Strategy/false-name proof protocols for combinatorial multi-attribute procurement auction: Handling arbitrary utility of the buyer. In: Proceedings of the First Workshop on Internet and Network Economics (WINE-2005). (2005)
14. Yokoo, M.: The characterization of strategy/false-name proof combinatorial auction protocols: Price-oriented, rationing-free protocol. In: Proceedings of the 18th International Joint Conference on Artificial Intelligence. (2003) 733–739
15. Moulin, H.: Proportional scheduling, split-proofness, and merge-proofness. Games and Economic Behavior
16. Du, Y., Sami, R., Shi, Y.: Path Auction Games When an Agent Can Own Multiple Edges. In: Proc. 1st Workshop on the Economics of Networked Systems (NetEcon06). (2006) 48–55
17. Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R.: Linear programming and vickrey auctions. IMA Volume in Mathematics and its Applications, Mathematics of the Internet: E-auction and Markets **127** (2001) 75–116
18. Watts, D.J., Strogatz, S.H.: Collective dynamics of 'small-world' networks. Nature **393**(6684) (June 1998) 440–442
19. Rochet, J.C.: A necessary and sufficient condition for rationalizability in a quasi-linear context. Journal of Mathematical Economics **16** (1987) 191–200