

Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions

Robert W. Day
Operations and Information Management
School of Business
University of Connecticut
Storrs, CT 06269-1041

S. Raghavan
Decision and Information Technologies
The Robert H. Smith School of Business
University of Maryland
College Park, MD 20742-1815

January 31, 2005

Abstract

Motivated by the increasing use of auctions by government agencies, we consider the problem of fairly pricing public goods in a combinatorial auction. A well-known problem with the incentive-compatible Vickrey-Clarke-Groves (VCG) auction mechanism is that the resulting prices may not be in the core. Loosely speaking, this means the payments of the winners could be so low, that there are losing bidders who would have been willing to pay more than the payments of the winning bidders. Clearly, this “unfair” outcome is unacceptable for a public-sector auction. Proxy-based combinatorial auctions, in which each bidder submits several package bids to a proxy, result in efficient outcomes and bidder-Pareto-optimal core-payments by winners, thus offering a viable practical alternative to address this problem.

This paper confronts two critical issues facing the proxy-auction. First, motivated to minimize a bidder’s ability to benefit through strategic manipulation (through collusive agreement or unilateral action), we demonstrate the strength of a mechanism that minimizes total payments among all possible proxy auction outcomes, narrowing the previously broad solution concept. Secondly, we address the computational difficulties of achieving these outcomes with a constraint-generation approach, promising to broaden the range of applications for which the proxy-auction achieves a comfortably rapid solution.

keywords: *combinatorial auctions, core allocations, bidder-Pareto-optimality, constraint generation, VCG payments, proxy auctions*

1 Introduction

Classic auction theory shows that if a single item is auctioned via the submission of sealed-bid price offers, then bidders can be expected to report their bids honestly under a second-price mechanism (i.e., one in which the highest bid wins and pays the amount of the second highest bid). Such a mechanism is said to be *incentive compatible*, meaning that truth-telling is a dominant strategy for every player, and that there is no incentive for unilateral deviation from the truth-telling strategy.

The analogous second-price mechanism for combinatorial auctions (in which bidders bid on combinations of items) is the well-known Vickrey-Clarke-Groves (VCG) mechanism, which has also been shown to be incentive compatible (for primary sources see Clarke, 1989; Groves, 1973; Vickrey, 1961). In the general version of this auction mechanism, bidders each submit a price for each possible combination of items.

Winners are chosen to maximize the combined social value of awarded bundles (with no item going to more than one bidder), and such an allocation is referred to as *efficient*. Each winner in the VCG auction then pays an amount less than her bid on the bundle she is awarded, receiving a discount from her actual bid that is calculated to eliminate her ability to gain from falsifying her preferences. In particular, each winning bidder receives a discount equal to the difference in value between the efficient solution with all bidders and the efficient solution in the absence of that particular bidder. The problem of finding an efficient allocation for a given set of bids is known as the *winner determination problem*, and is well known to be computational difficult (\mathcal{NP} -hard, see Rothkopf et al., 1998).

Though the VCG mechanism is widely discussed in the auction literature, its drawbacks are so numerous that it is rarely if ever used in practice. Indeed, there are too many problems to name here; we therefore direct the reader to a few good references on the problems with VCG mechanisms (see Ausubel and Milgrom, 2002; Rothkopf and Harstad, 1995; Rothkopf et al., 1990; Sakurai et al., 1999). Principle among these drawbacks is that the VCG payments may be so low (in a forward auction) that the outcome is not a “core” allocation (defined formally in the next section). Roughly stated, this means that a dissatisfied coalition of bidders may be able to suggest an auction outcome that is preferred by all the members of the coalition *and* the seller. This outcome is clearly unsatisfying, especially when the seller is a government agency assigned the duty of fair allocation, and the winners *have not paid enough* to establish themselves as the fair recipients of the auction items relative to the competition.

Given that the VCG mechanism is undesirable in practice, one may next ask, what other payment mechanism can be used in a combinatorial auction? The most obvious alternative is a first-price (or pay-as-bid) mechanism, but this too has drawbacks. In a sealed-bid combinatorial auction this payment rule encourages the bidders to submit bids that are just barely enough to achieve the efficient allocation. Indeed, bidding higher than the minimum amount necessary to secure a particular bundle simply gives more money to the seller in the pay-as-bid scenario. But if each player is trying to predict the minimum amount needed to win his efficiently awarded bundle, uncertainty about the bids of others makes this a precarious endeavor. With incomplete information about the preferences of others, trying to bid the minimum possible amount will often lead to bidding not quite enough, leading the auction mechanism to miss the efficient outcome as a consequence.

Unsatisfied with the two “extreme” possibilities for a payment mechanism (first-price and VCG), is there some other “second-price” combinatorial auction payment mechanism that maintains the desirable properties of the second-price mechanism for a single-item auction, while not suffering from the undesirable properties of the VCG mechanism? Most desirable among these second-price properties is strategic ease for the bidders. Specifically, in the second-price single-item auction and in the VCG auction, a bidder is not penalized for

bidding above the minimum amount necessary to win. Are there VCG-like mechanisms in which it is not necessarily a bad idea to bid straightforwardly or honestly (as it is in a first-price mechanism)?

Few such mechanisms have been studied extensively, and even fewer display these “second-price” properties while also maintaining efficiency, which we consider a prime motivation, indispensable given our interest in government allocation problems. At least one prominent example of such a mechanism is given by Ausubel and Milgrom (2002), whose *ascending proxy auctions* are fast becoming a viable practical alternative to the virtually unusable VCG mechanism for real-world combinatorial auction applications. Assuring competitive pricing and economic efficiency based on the bids, the proxy auction terminates at a desirable “bidder-Pareto-optimal core outcome” (defined formally in the next section). This presents an especially attractive solution for several public-good allocation problems, where the need for a “socially acceptable” outcome outweighs revenue maximization for the seller.

Designed as a proxy implementation of an iterative combinatorial auction, bidders in their mechanism are insulated against the dangers of bidding more than the minimum amount necessary to win a particular bundle, as the proxy submits bid information only at minimum increments. The particular algorithm for winner/payment determination in the ascending proxy auction (as first proposed by Ausubel and Milgrom, 2002) is interpreted by de Vries et al. (2003) as an implementation of the subgradient algorithm: at each moment in the process the algorithm solves an \mathcal{NP} -hard set-packing problem, determining a non-overlapping set of bundles from all those demanded by the bidders at some current set of prices, and a net utility maximizing bundle for each bidder. Each non-anonymous bundle price is then adjusted up by one increment for each bundle that is not contained in the “seller’s choice” allocation. This process corresponds to incremental stepping in the direction of a subgradient. As is commonly the case with the subgradient algorithm, convergence is slow and depends critically on the choices of step size.

Compounding the problem of rapid computation is that each iteration of this algorithm requires the solution of an integer program (IP) for winner determination. Although each IP can be solved reasonably quickly given advanced computational techniques (see, for example Günlük et al., 2005) and software (like CPLEX and XPRESS for example), the repeated solution and slow convergence inhibits practical implementation (see Hoffman et al., 2005).

Even since the relatively new development of the ascending proxy auction, a few methods for arriving at the proxy auction outcome more rapidly have been proposed, including the technique developed in this paper. Hoffman et al. (2005) give experimental evidence for slow performance of the subgradient-type ascending proxy implementation and show how to improve the computational speed of the proxy-auction’s iterative implementation by starting at the VCG solution, and by using a more sophisticated adjustment (scaling) of the price increment. Their technique requires that the auctioneer be allowed to “open up the proxy” and use

all of the information reported to the proxy to solve instances of the winner determination problem. Given that the auction mechanism can be trusted not to make the proxy information public (which may be legally enforced), we argue that such revelation is reasonable and should be implemented for large scale applications, since access to the proxy reports may greatly accelerate computational performance. Wurman et al. (2004) provide a different approach to accelerating the proxy implementation, by computing the “inflection” or “change” points in the iterative auction’s price trajectory. In §4, we look at these methods in greater depth and compare them to the technique presented here.

More recently, Ausubel et al. (2005) provide a practical *clock*-proxy auction design that terminates with bidders submitting combinatorial bids in a proxy auction after some initial rounds designed to reveal price information. The final proxy phase is strategically equivalent to a sealed-bid combinatorial auction, justifying our focus on sealed-bid price-mechanisms in the current treatment. Given the widespread applicability of their technique in public sector markets for electricity generation, spectrum licenses, and airport landing-slot rights (to name a few), the ability to produce final payments more rapidly and to better understand the selection of payment outcomes among the various possible core solutions may have a deep impact on our ability to successfully implement these market mechanisms. Indeed, researchers associated with the FCC and FAA have recently started experimenting with an implementation of the algorithm presented in this paper as a possible pricing mechanism for future auctions.

In this paper, we provide a new, more direct computational procedure for arriving at bidder-Pareto-optimal core outcomes, using constraint generation and an explicit definition of the core region in payment space. Our technique may be described as an *approximate VCG mechanism*, and we provide concrete justification for the use of these mechanisms in general. At the crux of our approach, we show how to overcome one of the pitfalls of a linear programming (LP) approach to payment determination in the core, using constraint generation to handle the exponential number of constraints necessary to define the core. In order to accomplish this, we formulate the *core separation problem*, finding the most violated core constraint (most upset coalition) for any proposed payment vector. This core separation problem is computationally equivalent to the winner determination problem, and we show how to integrate the separation technique into a procedure that settles on a bidder-Pareto-optimal core point, with no need for limiting (substitutability) assumptions on the preferences of the bidders. Along the way, we solidify the concept of a *coalitional contribution*, measuring a winning bidder’s desire to join a coalition at a particular payment vector. We demonstrate the flexibility of our technique by showing a few variations, and discuss the selection of a particular core outcome when several meet the “bidder-Pareto-optimal” criterion. In particular, we show that a mechanism that *minimizes total payments* consequently minimizes the total availability of gains from unilateral strategic manipulation, and is immune to a certain form of group collusion which would be

profitable in any mechanism without this property. Finally, we compare the computational performance of our technique to others in the literature with a few detailed examples, and discuss directions for future research.

We begin the next section with the introduction of our notation, followed by a motivating example and a theorem supporting the approximate VCG approach. In §3 we develop the algorithm itself, with a comparison to existing techniques in §4. These comparisons demonstrate the computational effectiveness of the algorithm for computing bidder-Pareto-optimal payments presented in §3, and reinforce our arguments on the *selection* of a bidder-Pareto-optimal payment vector from §2 by showing discrepancies among various solution techniques. Finally, we provide concluding remarks in §5.

2 Approximate VCG Mechanisms

Consider an environment where N distinct items are to be auctioned among M bidders. We will typically index each item in the set $I = \{1, 2, \dots, N\}$ of auction items by the letter i , and each bidder in the set $J = \{1, 2, \dots, M\}$ of all bidders by the letter j . For any set $S \subseteq I$ let $v_j(S)$ denote bidder j 's value for the bundle of items S (the maximum amount he would be willing to pay for S), and let $b_j(S)$ denote the bid that bidder j has submitted to the auction for the bundle S . To maintain full generality, we assume an *XOR* language throughout (i.e., every bidder submits an exclusive bid for every possible bundle S), though the technique we propose in §3 easily generalizes to any bidding language that uses an IP formulation for winner-determination. Further we will use b_j and v_j to denote the full reports and valuations of bidder j over all bundles $S \subseteq I$.

The variable π will be used to denote a payment vector, with each component π_j indicating the payment made by bidder j , and with superscripts used to distinguish among different payment vectors. For example, π_j^{VCG} will indicate bidder j 's VCG payment, while π^t will indicate the vector of payments for all bidders at iteration t in our payment adjustment procedure. In Theorem 2.1, notation of the form $\pi(b_j, b_{-j})$ will be used to denote the payment vector determined after a report of b_j by bidder j and a fixed set of bids b_{-j} by her competitors.

An *outcome*, Γ , refers to an allocation and set of payments for bidders in a combinatorial auction. Let the *coalition*, C_Γ , refer to the set of bidders receiving items under outcome Γ .

Perhaps the most obvious problem with VCG payments in terms of ex-post satisfaction is that the VCG outcome may not be a “core outcome”. The core conditions for a one-sided (forward) auction can be stated by the following definitions, where a bidder *weakly prefers* outcome Γ_1 to Γ_2 if Γ_1 gives him utility greater than or equal to the utility of outcome Γ_2 . We assume quasilinear net utility throughout (i.e. utility of a

bundle S to bidder j , is simply $v_j(S) - \pi_j$.

Definition: An outcome Γ is *blocked* if there is an alternative outcome Γ_B which generates more revenue for the seller and for which every bidder in C_{Γ_B} weakly prefers Γ_B to Γ . C_{Γ_B} may be referred to as a *blocking coalition*.

Definition: An outcome Γ that is not blocked is called a *core outcome*.

Definition: A core outcome is *bidder-Pareto-optimal* if there is no other core outcome weakly preferred by every winning bidder

We take as our primary motivation the well-known fact that a VCG outcome may not be a core outcome, and thus may be considered “socially unacceptable,” and unfit for an auction of public goods. Consider the following three-bidder, two-item example from Ausubel (2000): Let $b_1(\{A, B\}) = b_2(\{A\}) = b_3(\{B\}) = 2$. In the efficient allocation, bidder 2 wins item A while bidder 3 wins item B . VCG payments for each winning bidder are computed to be zero. This is not a core outcome because bidder 1 would prefer to pay any amount up to 2 units to receive both items, an outcome which is clearly more desirable to the seller than VCG outcome, in which no payments are collected. Even worse, if bidders 2 and 3 value the items at $v_2(\{A\}) = v_3(\{B\}) = 0.5$, they can still place bids of 2 units each, and under the VCG mechanism will be not be held financially accountable for displacing the efficient winner, bidder 1. If instead bidders 2 and 3 each pay 1 given the bids of 2 units each, the coalition containing just bidder 1 no longer blocks, and the outcome is in the core. Further, the joint deviation from truthful bidding that benefitted bidders 2 and 3 when their valuations were each less than 1 is no longer profitable. Given this seemingly fatal flaw of the VCG mechanism, it is important in a public sector auction to eschew VCG payments in favor of a more reasonable payment determination mechanism that will be guaranteed to arrive at a core outcome.

Breaking away from the VCG mechanism is often met with skepticism from auction theorists: what strategic behavior can we expect from bidders in a non-incentive-compatible auction? Theorem 2.1 provides a bound on a bidder’s ability to benefit by deviating from the truthful reporting strategy. In particular, the difference between a bidder’s payment using the VCG mechanism and an alternative mechanism (assuming truthful reports) is the maximum that a bidder can gain by deviating in the alternative mechanism. Using the corollary that follows, these bounds provide us with the ability to minimize the total gains from deviation (over all bidders). These results provide support for the approximate VCG approach, answering the aforementioned skepticism in the most intuitive fashion: the closer you get to the VCG payments the less incentive there is to deviate.

Suppose that each bidder j reports a vector b_j , and that instead of assigning VCG payments $\pi^{VCG}(b_j, b_{-j})$ our auction mechanism uses a payment rule $\pi(b_j, b_{-j})$ that is guaranteed to assign payments greater than or equal to the VCG payments. (This assumption is justified because core payments are always greater than

or equal to the VCG payments.) Also, assume that the mechanism chooses the efficient allocation (based on the reports), so that for the same reports the mechanism in question differs from the VCG mechanism in payments only (i.e., allocations are the same under both mechanisms for the same reports). We then have the following result:

Theorem 2.1. *The amount that bidder j can benefit by unilaterally deviating from the honest report strategy is less than or equal to $\pi_j(v_j, b_{-j}) - \pi_j^{VCG}(v_j, b_{-j})$.*

Proof. Suppose not: There is some report \hat{b}_j such that:

$$(v_j(\hat{S}_j) - \pi_j(\hat{b}_j, b_{-j})) - (v_j(S_j) - \pi_j(v_j, b_{-j})) > \pi_j(v_j, b_{-j}) - \pi_j^{VCG}(v_j, b_{-j})$$

where \hat{S}_j is the bundle awarded to bidder j given the reports (\hat{b}_j, b_{-j}) , and S_j is the efficiently awarded bundle given the report (v_j, b_{-j}) . After rearranging and canceling, we have

$$v_j(\hat{S}_j) - \pi_j(\hat{b}_j, b_{-j}) > v_j(S_j) - \pi_j^{VCG}(v_j, b_{-j}).$$

By assumption on the mechanism π (and a standard assumption of quasilinear utility) we have

$$v_j(\hat{S}_j) - \pi_j^{VCG}(\hat{b}_j, b_{-j}) \geq v_j(\hat{S}_j) - \pi_j(\hat{b}_j, b_{-j}),$$

and thus from the above two inequalities we find

$$v_j(\hat{S}_j) - \pi_j^{VCG}(\hat{b}_j, b_{-j}) > v_j(S_j) - \pi_j^{VCG}(v_j, b_{-j}).$$

But this contradicts the well-known incentive-compatibility property of the VCG mechanism, thus the supposition must be false. \square

This theorem and its elementary proof lend theoretical support to several payment setting techniques for combinatorial auctions (or two-sided exchanges) which eschew the VCG mechanism, but approximate the VCG outcome as closely as possible within a feasible set of payments. These bounds are first stated in Parkes et al. (2001), who also demonstrate that they are tight. We provide a more rigorous proof of these bounds and use them to justify the approach of approximating VCG payments for one-sided auctions.

For example, suppose we want to determine a payment rule that assigns payments to winning bidders in an efficient allocation, and additionally that these payments must satisfy some prespecified properties (for example, the payments are not blocked by any coalition of bidders). If R is the region in payment space

containing all payment vectors satisfying the prespecified properties (for example the core), assuming that R is bounded below by the VCG payments, we have the following result:

Corollary 2.2. *The payment rule that minimizes total potential gains from deviation within the feasible region R minimizes total payments over R .*

Proof. When minimizing the sum of the (tight) bounds provided by Theorem 2.1, the VCG payment terms are constant relative to the payments, with the result following. \square

In Parkes et al. (2001), VCG payments were approximated subject to budget-balance in a two-sided combinatorial exchange. The authors investigated several payment rules for VCG approximation, eventually favoring a “threshold rule” in which payments are selected that minimize the maximum difference from VCG payments over all players. In the case of a budget-balanced exchange, however, total payments are constant (equal to zero), and thus Corollary 2.2 has no bite. In the case of a one-sided auction in which VCG payments are approximated within the core, one approach may be to simply follow the lead of Parkes et al. (2001) and attempt to find payments that minimize the maximum difference from VCG payments within the core. Corollary 2.2, however, shows that we must first minimize total payments, or else the global ability to gain from deviation may not be minimized. In §4 we provide an example demonstrating that a set of payments that minimizes the maximum difference from the VCG payments may not minimize total payments, and our algorithm developed in §3 shows how to incorporate the minimization of the maximum deviation from VCG as a secondary objective.

Having established that our goal is to minimize total payments within the core region in payment space, we now explore the problem of properly defining the core using linear and integer programming (IP). In order to do so, we will make use of an IP formulation for a General version of the Winner Determination problem:

$$\begin{aligned} & \max \sum_{j \in J} \sum_{S \subseteq I} b_j(S) \cdot x_j(S) && \text{(GWD)} \\ \text{subject to} & \sum_{S \supseteq \{i\}} \sum_{j \in J} x_j(S) \leq 1, && \forall i \in I && (1) \\ & \sum_{S \subseteq I} x_j(S) \leq 1, && \forall j \in J && (2) \\ & x_j(S) \in \{0, 1\}, && \forall S \subseteq I, \forall j \in J && (3) \end{aligned}$$

The solution of this problem finds the efficient allocation for a given set of bids, represented by the coefficients $b_j(S)$ in the objective function. In this formulation, binary variables $x_j(S)$ equal 1 if and only if bidder j

is awarded bundle $S \subseteq I$. (Constraints (3) tell us that these variables must be binary). Constraints (1) consequently ensure that each item is assigned to at most one bidder, while constraint set (2) prevents the auctioneer from accepting multiple bids from the same bidder. With “demand” constraints formulated as in (2), each bid is an exclusive offer made by a bidder for a particular bundle, and may not be combined or accepted in conjunction with any other offer from that bidder. The implied *bidding language* is therefore often referred to as an *XOR* language.

In the following section we provide a novel contribution, an algorithm for determining core payments using constraint generation. Our technique requires the repeated use of a slightly modified winner determination problem, for which GWD will serve as a starting point. Though the *XOR* language may seem cumbersome to those who appreciate the benefits of a more sophisticated bidding language, we present our results in the *XOR* language because it is fully expressive, allowing for the expression of any demand function through complete enumeration, and therefore most general. The method we present, however, may be applied in the context of *any* bid language, by simply performing the technique introduced in the following section on the corresponding winner determination problem for any particular bidding language.

3 Core Constraint Generation

As in Hoffman et al. (2005), our general approach is to first solve several winner determination problems for a sealed-bid combinatorial auction, settling on a particular set of winning bidders and VCG payments. Given this fixed outcome Γ_{VCG} we denote the associated coalition of winners as $W = C_{\Gamma_{VCG}}$ and the associated payments π_j^{VCG} for each bidder $j \in W$. Starting from this VCG outcome, we wish to arrive at a bidder-Pareto-optimal core outcome with the same efficient allocation; for the same bids only payments will differ between the VCG outcome and the final core outcome proposed here.

From a constrained optimization point of view it is difficult to categorize the bidder-Pareto-optimal core payments for two reasons. First, an exponential number of possible coalitions must be considered to define the core region in payment space; each coalition (subset of the set J of all bidders) is willing to make some offer to the seller, and the total payments of the winning bidders must exceed each of these offers. Secondly, it is difficult to gauge how much a winning bidder will contribute to a “coalitional value function” (the offer made to the seller) without knowledge of his final payment.

Given the core definition provided in §2, it seems convenient to define a core constraint for any blocking outcome Γ as $\sum_{j \in W} \pi_j \geq z_C$, where π_j is a payment made by each bidder $j \in W$, and z_C is the *coalitional value* of C . We define this coalitional value z_C as the maximum total payments that a coalition C would be willing to offer the seller in any outcome Γ with $C_\Gamma = C$. This formulation of a core constraint emphasizes

that any offer that could be made by a coalition must be matched by the set of winning bidders. With every coalition defining a constraint of this form, we suggest our first characterization of the core:

$$\sum_{j \in W} \pi_j \geq z_C, \quad \forall C \subseteq J \quad (\text{CORE1})$$

$$\pi_j^{VCG} \leq \pi_j \leq b_j(S_j), \quad \forall j \in W$$

where the VCG payments can be used as lower bounds on core payments (if they have been computed), but are not necessary to define the core. Additionally, we make use of the bid amounts themselves as upper bounds on payments in our LP formulation, since a bidder forced to pay more than his bid would dissent, preferring not participate in the auction.

Defining the core points in payment space with linear constraints gets us closer to computing the bidder-Pareto-optimal core payments using an LP, and allows us to restate our two difficulties from above:

- Noting the exponential number of constraints in CORE1, how do we separate which constraints must be applied from those that can be ignored without consequence?
- How do we compute the value of z_C for a particular coalition C ?

To appreciate the subtlety of the second question, observe that any bidder $j \in W$ would not want to join a coalition and receive less surplus than he would at his current payment π_j^t , at any iteration of a price-adjustment algorithm or point in time t . Trying to solve the winner determination problem restricted to the bidders in the coalition C therefore overstates the amount that any winning bidder in C would contribute to the coalitional value z_C , because he will not be compensated his opportunity cost (the amount of surplus he stands to gain at the current vector of payments). The result would be the application of a too restrictive constraint, therefore charging the winning bidders too much. We must instead consider that each winning bidder has an opportunity cost limiting his contributions to a blocking coalition. We therefore amend our notation to capture the fact that a coalitional value function depends on the current set of payments π^t , and thus change to $z_C(\pi^t)$.

We may now describe our general iterative approach as follows: at a current vector in payment space, π^t at each discrete iteration t , find the coalition C^t with the highest coalitional value relative to current payments, $z_C(\pi^t)$. If this coalition blocks the currently proposed allocation (the efficient solution with payments π^t), apply the corresponding core constraint and find a new bidder-Pareto-optimal payment π^{t+1} . If this coalition does not block then terminate at the currently proposed outcome. Before demonstrating this technique in full detail, we observe a few basic properties of each winning bidder's "coalitional contribution"

relative to current payments, leading us to a more useful LP characterization than that of CORE1.

Let S_j be the bundle awarded to winning bidder j in the efficient solution. Relative to the current payment vector, bidder j would not voluntarily join a coalition that offers him less surplus than $b_j(S_j) - \pi_j^t$, his opportunity cost. If a coalition provides bidder j with bundle \bar{S}_j (possibly the same as S_j , but in general not) then bidder j would contribute at most $b_j(\bar{S}_j) - (b_j(S_j) - \pi_j^t)$ into the coalition value function; if he contributed more, then he would perceive less benefit (surplus) from the hypothetical outcome of the coalition than the one proposed by the auctioneer at iteration t . If bidder j is not in the set W , there is no opportunity cost to recover, and he would offer his entire value for a bundle in an effort to block the set of winning payments. In this case $b_j(S_j = \emptyset) = \pi_j^t = 0$ and his coalitional contribution would be $b_j(\bar{S}_j)$. In general we have the following definition:

Definition: Facing payment of π_j^t for bundle S_j in the efficient allocation, bidder j would be willing to make a *coalitional contribution* of $q_j(C, \pi_j^t) = b_j(\bar{S}_j) - b_j(S_j) + \pi_j^t$ to receive \bar{S}_j as part of the coalition C .

Lemma 3.1. *We note the following basic properties of the coalitional contribution $q_j(C, \pi_j^t)$:*

1. *If bidder j were to pay an amount π_j^C that is greater than $q_j(C, \pi_j^t)$ to join coalition C and win \bar{S}_j , then bidder j would experience less surplus from the coalition than from the auctioneer's proposed outcome S_j with payment π_j^t*
2. *$q_j(C, \pi_j^t)$ increases linearly in π_j^t , $\forall j \in W$*
3. *$\forall j \in J$, $\bar{S}_j = S_j$ implies that $q_j(C, \pi_j^t) = \pi_j^t$*
4. *$q_j(C, \pi_j^t) = b_j(\bar{S}_j)$ and $j \in W \cap C$ imply that $\pi_j^t = b_j(S_j)$*
5. *$q_j(C, \pi_j^t) = b_j(\bar{S}_j)$, $\forall j \notin W$*

Proof. Property (1) verifies that our definition of the quantity $q_j(C, \pi_j^t)$ truly reflects what we mean by coalitional contribution, and follows from the definition: $\pi_j^C > q_j(C, \pi_j^t) = b_j(\bar{S}_j) - b_j(S_j) + \pi_j^t$ implies that surplus from the coalitional outcome is $b_j(\bar{S}_j) - \pi_j^C < b_j(S_j) - \pi_j^t$, where the right-hand-side of the last inequality is exactly the surplus from the auctioneer's outcome. Properties (2), (3), and (4) follow directly from the definition while property (5) follows from the standard assumptions that $b_j(\emptyset) = 0 \forall j$, and that $\pi_j^t = 0 \forall j \notin W$. □

Now that we have established the importance of opportunity cost and shown that bidder j 's coalitional contribution will normally be less than his value for the bundle given to him by the coalition, we formulate the *core-constraint separation problem* at payment vector π^t . At any point in payment space, π^t , the integer

program SEP^t finds the most violated core constraint, or tells us that no blocking coalition can be found.

$$\begin{aligned}
z(\pi^t) &= \max \sum_{j \in JS \subseteq I} \sum_{S \subseteq I} b_j(S) \cdot x_j(S) - \sum_{j \in W} (b_j(S_j) - \pi_j^t) \cdot \gamma_j & (\text{SEP}^t) \\
\text{subject to} \quad & \sum_{S \supseteq \{i\}} \sum_{j \in J} x_j(S) \leq 1, \quad \forall i \in I \\
& \sum_{S \subseteq I} x_j(S) \leq 1, \quad \forall j \in J \setminus W \\
& \sum_{S \subseteq I} x_j(S) \leq \gamma_j, \quad \forall j \in W \\
& x_j(S) \in \{0, 1\}, \quad \forall S \subseteq I, \forall j \in J \\
& \gamma_j \in \{0, 1\}, \quad \forall j
\end{aligned}$$

The added terms in the objective of this IP formulation (starting from GWD from §2) ensure that any winning bidder will be compensated his opportunity cost if selected as part of the optimal solution to SEP^t . It is easy to verify that any bidder $j \in J$ contributes exactly $q_j(C, \pi_j^t)$, his coalitional contribution, and that this formulation is equivalent to an instance of the winner determination problem for the auction in which each bid from a winning bidder is reduced by his current opportunity cost.

If the objective $z(\pi^t) > \sum_{j \in W} \pi_j^t$ then the bidders for which $x_j(S) = 1$ form a coalition C^t which blocks the efficient allocation with the current set of payments; i.e. C^t is a set of bidders such that $z(\pi^t) = z_{C^t}(\pi^t)$. If $z(\pi^t) = \sum_{j \in W} \pi_j^t$ then we have achieved an unblocked core outcome; no coalition of bidders would be able to offer an outcome that both they and the auctioneer would prefer. Also note that $z(\pi^t) < \sum_{j \in W} \pi_j^t$ is not a possibility since the feasible allocation of items to the winning set W achieves an objective value of $\sum_{j \in W} \pi_j^t$. As a result, the algorithm developed in this paper is *ascending* in terms of total payments.

Having found a coalition C^t that blocks the proposed outcome with payment vector π^t , we know that the constraint $\sum_{j \in W} \pi_j \geq z_{C^t}(\pi^t)$ is violated. One way to characterize the set of core payments is therefore as the set of points in payment space that satisfy the constraints formulated above as CORE1. The region defined by CORE1, however, is difficult to generate or optimize over because there are an exponential number of constraints in the first set and each has a right-hand-side requiring optimization of an \mathcal{NP} -hard winner determination problem to compute. In addition, property (2) from Lemma 3.1 indicates that price increases on any bidder $j \in W \cap C$ completely cancel with the corresponding increase in coalitional contribution $q_j(C, \pi_j^t)$ on the right-hand-side of the constraint for C in CORE1.

Consequently, our method is to find bidder-Pareto-optimal core payments by generating core constraints of the form $\sum_{j \in W \setminus C^t} \pi_j \geq z(\pi^t) - \sum_{j \in W \cap C^t} \pi_j^t$. The right-hand-side of these constraints are computed

once and *remain constant for the remainder of the algorithm*, unlike the right-hand-side of the constraints in CORE1. In economic terms, we choose these constraints because they do not hold winning bidders in a blocking coalition C^t responsible for overcoming coalition C^t , the futility of which is clear by our formulation of coalitional contribution.

After finding each of these constraints using SEP^t, we may solve the following linear program BPO^t to find a set of bidder-Pareto-optimal core payments, in the core relative to all coalitions found through iteration t . Our choice of objective function is motivated by Corollary 2.2.

$$\theta^t = \min \sum_{j \in W} \pi_j \quad (\text{BPO}^t)$$

$$\begin{aligned} \text{subject to} \quad & \sum_{j \in W \setminus C^\tau} \pi_j \geq z(\pi^\tau) - \sum_{j \in W \cap C^\tau} \pi_j^\tau, \quad \forall \tau \leq t & (\text{CORE2}) \\ & \pi_j^{VCG} \leq \pi_j \leq b_j(S_j), \quad \forall j \in W \end{aligned}$$

We then use the value of each π_j in the solution for the next iteration (i.e. set $\pi_j^{t+1} = \pi_j$). We now show that the resulting ascending algorithm converges to the desired outcome.

Theorem 3.2. *If $z(\pi^t) = \theta^{t-1}$, then the solution to BPO^{t-1} yields bidder-Pareto-optimal core payments.*

Proof. First we note that the payments π^t are in the core, because if not a blocking coalition provides a solution to SEP^t greater than θ^{t-1} . By the minimality of $\sum_{j \in W} \pi_j$ with respect to the constraints (CORE2), if any bidder experiences a payment reduction without an increase in the payment of some other bidder, then some constraint from (CORE2) must be violated. Thus, we find that the payments must be Pareto-efficient with respect to core constraints. The convergence of this algorithm is guaranteed because only a finite number of constraints may be generated, and because the region (CORE2) always contains at least the trivial pay-as-bid solution. \square

In some cases, there may be multiple optimal solutions to BPO^t, for which we suggest the following

Table 1: The Core Constraint Generation Algorithm

| |
|--|
| <p>Initialize: Iteratively solve Winner Determination Problem, finding winners W and VCG payments; Set $\pi^1 = \pi^{VCG}$, and $\theta^0(\varepsilon) = \sum_{j \in W} \pi_j^{VCG}$</p> |
| <p>Iteration t: Solve SEP^{t} If $z(\pi^t) > \theta^{t-1}(\varepsilon)$ Add constraint $\sum_{j \in W \setminus C^t} \pi_j \geq z(\pi^t) - \sum_{j \in W \cap C^t} \pi_j^t$ to EBPO^{t} and solve, updating $\theta^t(\varepsilon)$. Set $\pi_j^{t+1} = \pi_j$, the values found from EBPO^{t}. Iterate: $t = t + 1$ Else Terminate</p> |

refinement, a linear program EBPO ^{t} which finds *equitable* bidder-Pareto-optimal core payments:

$$\theta^t(\varepsilon) = \min \sum_{j \in W} \pi_j + \varepsilon m \tag{EBPO ^{t} }$$

subject to

$$\sum_{j \in W \setminus C^\tau} \pi_j \geq z(\pi^\tau) - \sum_{j \in W \cap C^\tau} \pi_j^\tau, \quad \forall \tau \leq t$$

$$\pi_j - m \leq \pi_j^{VCG}, \quad \forall j \in W$$

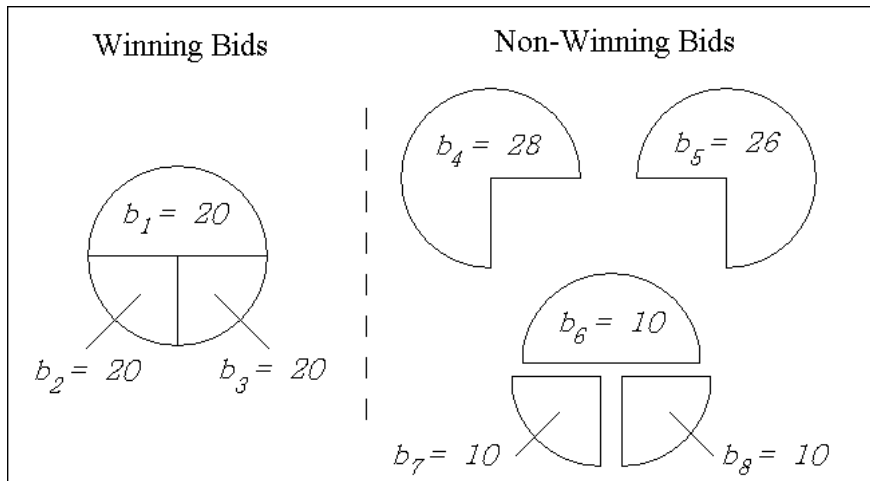
$$\pi_j \leq b_j(S_j), \quad \forall j \in W$$

$$\pi_j \geq \pi_j^{VCG}, \quad \forall j \in W$$

with decision variables appearing only on the left-hand-side of each constraint. By taking a small enough value of the scalar ε , we have that $\sum_{j \in W} \pi_j$ computed from the solution to this problem is equal to θ^t from the corresponding instance of BPO ^{t} . We then see that the effect of the new terms is to find a set of payments that minimizes the maximum difference from the VCG payments over all bidders, amongst all outcomes that minimize the total payments of winning bidders. The entire process for determining equitable bidder-Pareto-optimal core payments is summarized in Table 1.

Parkes (2002) discusses several core outcomes related to earlier payment selection mechanisms for two-sided auctions with budget-balance in Parkes et al. (2001). Among these are a minimax deviation from the VCG payments rule (or threshold rule), and a maximin deviation from the bundle values (or equal-pay rule). The modified objective function in EBPO ^{t} allows us to mimic the equity properties of the Parkes et al. (2001) threshold rule as a secondary objective, while minimizing the total ability of bidders to gain

Figure 1: An Auction for which threshold payments do not minimize total payments



from deviation according to Corollary 2.2.

In the next section we illustrate this Core Constraint Generation algorithm with an example, and compare the results to those of the threshold rule for one-sided auctions as discussed by Parkes (2002), and to the examples presented in Hoffman et al. (2005) and Wurman et al. (2004). In §4.1, we establish the primacy of total-payment minimization over the space of all bidder-Pareto optimal solutions, and consider a related Nash Bargaining game in §4.2.

4 Examples and Comparisons to Other Techniques

To illustrate our technique for determining bidder-Pareto-optimal core outcomes using Core Constraint Generation, consider the three-item, eight-bidder example of Figure 1. Here we depict the three items as pieces of a pie to indicate visually which bids are on which items. Inspection shows that bidders 1, 2, and 3 constitute the set of winners in the efficient allocation, as indicated. In accordance with our earlier notation we will say that $W = \{1, 2, 3\}$.

With the set of winners determined, we next compute $\pi_1^{VCG} = \pi_2^{VCG} = \pi_3^{VCG} = 10$, and so set $\pi_1^1 = \pi_2^1 = \pi_3^1 = 10$. Solving SEP¹ the separation problem at iteration 1, we find the most violated blocking coalition $C^1 = \{3, 4\}$ with $z(\pi^1) = 38 > 30 = \sum_{j \in W} \pi_j^1$. We then add the constraint $\pi_1 + \pi_2 \geq 28$ to the formulation EBPO¹ and solve to find the new set of payments $\pi_1^2 = \pi_2^2 = 14$, $\pi_3^2 = 10$. Next we solve SEP² and find that the coalition $C^2 = \{2, 5\}$ blocks the current set of payments with $z(\pi^2) = 40 > 38 = \sum_{j \in W} \pi_j^2$. We form EBPO² by adding the constraint $\pi_1 + \pi_3 \geq 26$ and solve, yielding $\pi_1^3 = 16$, $\pi_2^3 = 12$, and $\pi_3^3 = 10$. Finally, solving SEP³ we find that no blocking coalitions exist; the process has terminated at a bidder-Pareto-optimal

core outcome.

4.1 The Threshold Rule and a Collusion Problem

The algorithm we have just demonstrated selects a specific Pareto-efficient core outcome. Elsewhere in the recent literature, different methods of selecting core outcomes are suggested (Parkes, 2002), at times resulting in different outcomes than the ones generated by our technique. We compare our method now to the “mini-max” rules for finding Pareto-efficient core outcomes that are analogous to the rules developed for finding balanced-budget solutions in combinatorial exchanges (Parkes et al., 2001). For example, reformulated in our own notation, the *threshold rule* selects payments which solve the following optimization problem:

$$\min \max_{j \in W} (\pi_j - \pi_j^{VCG}) \quad (\text{THRESH})$$

$$\text{subject to } \sum_{j \in W \setminus C} \pi_j \geq n_C, \forall C \subseteq J$$

where n_C denotes the portion of the coalitional value for C that is attributable to non-winners. This rule for payment selection can be computed as an LP:

$$\min m \quad (\text{THRESH-LP})$$

$$\text{subject to } \sum_{j \in W \setminus C} \pi_j \geq n_C, \forall C \subseteq J$$

$$\pi_j - \pi_j^{VCG} \leq m, \forall j \in W$$

Solving this LP for the above example, we find the solution $\pi_1 = \pi_2 = \pi_3 = 14$, which is unfortunately Pareto-inefficient. There is however a Pareto-efficient point amongst the optimal solutions to THRESH-LP, and we may next consider the strength of the threshold rule as a selection criteria amongst the multitude of points on the Pareto-frontier of the core. As in our own formulation for selecting equitable payments, this may be accomplished with the insertion of a secondary objective weighted by a tiny value ε :

$$\min m + \varepsilon \cdot \sum_{j \in W} \pi_j \quad (\text{Pareto-THRESH})$$

$$\text{subject to } \sum_{j \in W \setminus C} \pi_j \geq n_C, \forall C \subseteq J$$

$$\pi_j - \pi_j^{VCG} \leq m, \forall j \in W$$

Solving this new LP for the example of this section, we arrive at the solution $\pi_1 = \pi_2 = 14$, and $\pi_3 = 12$ which is indeed Pareto efficient and minimizes the maximum deviation from the VCG payments. We note, however, that in accordance with the bounds presented in Theorem 2.1, this solution provides 10 units of opportunity to gain from deviation, while the EBPO^t solution provides only 8 units of opportunity to gain from deviation, a global decrease in the opportunities from deviation, although individual bidders may have greater opportunities.

The selection of a Pareto-efficient outcome is by its very nature a matter of taste. Pareto-efficient points are by definition unable to be compared with a strict dominance relationship; movement from one Pareto optimal point to another results in an increase in utility to one player if and only if another player experiences a decrease in utility. Though it seems reasonable to guarantee that total available incentive to deviate be minimized, one could argue that minimizing the maximum incentive to deviate is more important (the previous example demonstrates that there is a distinction). We note, however, the following observation:

Proposition 4.1. *A core outcome that minimizes the total payments by bidders will be strictly preferred by the bidders to any other Pareto-efficient core outcome if side payments are possible.*

Since Ausubel and Milgrom (2002) show that collectively bidding any bidder-Pareto-optimal point in the core constitutes a Nash equilibrium, one must look outside the concept of unilateral deviation in order to compare bidder-Pareto-optimal outcomes. In order to select a payment mechanism from all bidder-Pareto-optimal choices, we must therefore consider multilateral collusion supported by side payments.

Compare, for example, the outcome arrived at by the solution of EBPO^t, $\pi_1 = 16$, $\pi_2 = 12$, $\pi_3 = 10$, and the outcome arrived at by the solution of Pareto-THRESH, $\pi_1 = \pi_2 = 14$, $\pi_3 = 12$. In the former case the sum of the payments equals 38, while in the latter case the payments sum to 40. If asked which outcome is preferred, in the presence of side payments the winners will select the former. This is because the bidders that stand to benefit from the change (bidders 2 and 3, in this case) can pay off the bidders who would experience a payment increase (just bidder 1) to make them indifferent to the change. In this case, bidders 2 and 3 could compensate bidder 1 a payment of 1 unit each in an effort to coax him into a move from the latter to the former. Bidder 1 would then be indifferent while bidders 2 and 3 will each still gain 1 unit of surplus from the change. We note that the same comparison holds for the “equal-pay rule” discussed by Parkes (2002), (in which the maximum difference between the bundle values and payments is minimized), since by the construction of the example in Figure 1 all winning bidders have the same valuations and VCG payments.

Proposition 4.2. *For any auction mechanism terminating at a bidder-Pareto-optimal core outcome, it is possible for the winning bidders to collude to bid (and pay) exactly the final payments prescribed under a*

total-payment-minimization mechanism (e.g., EBPO^t).

This result is achieved by using covert side payments to support any difference from the mechanism selected by the auctioneer, as shown in the preceding example. Further, as in the preceding example this collusion will be advantageous since total payments are reduced, leaving money available to support the covert side payments.

Interestingly, if the public were made aware of the side payments, and each bidder's apparent final payment for her awarded bundle is adjusted accordingly, the resulting solution is no longer in the core. For example, if bidders facing the threshold rule make side payments to support a move to the EBPO^t solution for the Figure 1 auction as mentioned earlier, the adjusted final payments (adding or subtracting side payments) are $\pi_1 = 14$, $\pi_2 = 13$, $\pi_3 = 11$. But this outcome is blocked; bidder 4 would be willing to pay up to 28 monetary units for the items won by bidders 1 and 2 for just 27 monetary units.

Since this possibility of colluding to a non-core point exists whenever the auction mechanism does not minimize total payments within the core, we argue that the auction mechanism should necessarily minimize total payments within the core. By selecting an outcome that would be emulated under the collusive strategy, we eliminate the incentive to engage in such collusion with illicit payments.

We conclude that there is a phenomenon present in the example of Figure 1 demonstrating a problem with the threshold rule as previously formulated for use in this context. For the original formulation of the threshold rule (with respect to budget-balance in a combinatorial exchange, as in Parkes et al., 2001) total payments are held fixed, explaining why the problems of a non-payment-minimizing mechanism might at first be overlooked when applying the threshold concept to the new realm of core solutions in a one-sided auction. Contrary to the formulation given by Parkes (2002), we have demonstrated strong arguments in favor of a payment rule that explicitly minimizes total payments in the context of one-sided auctions and core constraints.

4.2 Nash Bargaining

With arguments in the previous subsection proposing the selection of an outcome that winners may try to achieve through collusion under a different mechanism, we must naturally consider the connections to the theory of bargaining. In a seminal paper, Nash (1950) introduces a solution concept for a set of players bargaining over a convex set of feasible outcomes in payoff space. Under a few mild assumptions, it can be shown that the players should immediately settle on the point that maximizes the product of their utilities within this convex set. Since the convexity of the core is a well-known property, one may wonder if there is an interpretation of our auction mechanism as a Nash bargaining game, and whether the outcomes coincide.

After all, we are proposing to select the outcome that bidders can profitably collude to under some other mechanism. Is this the same solution that Nash bargaining says these bidders will naturally agree to? Consider the following game:

Bidders are told that they are participating in a VCG auction and therefore respond with their true valuation for every bundle. The auctioneer releases the auction (VCG) outcome including the efficient winners and their payments, causing great discontent among losing bidders because the solution is not in the core. As a political solution to the public outcry, the auctioneer asks the winning bidders to decide among themselves on a set of payments that no group of dissatisfied bidders could complain about (i.e., a core outcome), or else the auction results will be canceled with no trade. This is a Nash bargaining game with the core as a feasible region. The solution is the point in the core that maximizes the product of the winning bidder utilities.

Is the solution of this bargaining game the same as the solution found using Core Constraint Generation? For some examples yes, but in general the answer is no. Again, take the example in Figure 1. The unique solution to the Nash bargaining problem for this auction is $\pi_1 = 16 - \lambda$, $\pi_2 = 12 + \lambda$, $\pi_3 = 10 + \lambda$, where

$$\lambda = \frac{14 - 2\sqrt{43}}{3}$$

Finding this solution requires the solution of a non-linear programming problem, which in general may be quite difficult. It is easy to verify that the product of all winning bidders' utilities is higher under this solution than the one found using Core Constraint Generation. To verify that this is indeed the optimal solution to the Nash bargaining problem one need only verify the pseudoconcavity of the objective function (the product of three quasilinear non-negative utility functions) over the core, and that this point satisfies the Karush-Kuhn-Tucker optimality conditions (see, for example Bazaraa et al., 1979). The derivation of this solution is a bit tedious, but would make a healthy exercise for a student of non-linear programming.

Is this a better core solution for this auction than the one found using Core Constraint Generation? We first note that it suffers from the same problem as the threshold solution (noted in §4.1) that total payments are not minimized (payments sum to $38 + \lambda$, rather than the minimal 38). Thus if an auction mechanism selects the payment vector within the core that maximizes the product of apparent utility, there will be a greater total incentive to deviate via Corollary 2.2. Further, taking into consideration the comments on side payments from §4.1, if the winning bidders were allowed to select a “socially acceptable” outcome within the core as depicted above, they would prefer to choose a payment minimizing outcome, with bidders 2 and 3 each covertly paying bidder 1 the amount $\lambda/2$. The bidders would announce the “socially acceptable” outcome $\pi_1 = 16$, $\pi_2 = 12$, $\pi_3 = 10$, but with side payments would be achieving the “socially unacceptable”

outcome $\pi_1 = 16 - \lambda$, $\pi_2 = 12 + \frac{\lambda}{2}$, $\pi_3 = 10 + \frac{\lambda}{2}$. The Nash bargaining model does not take into consideration the expansion of the feasible region by use of side payments, nor does it consider the incentive properties of the auction giving rise to this feasible region, though it is interesting to note the connection and the disparity between the two solution concepts.

We conclude by noting that if we amend the notion of the set of “socially acceptable” outcomes in the Nash game to be the set of total-payment-minimizing points in the core, than we get a new interesting choice for a “secondary objective.” The set of total-payment-minimizing points in payment space form a hyper-plane of one less dimension than the payment space itself, and finding the Nash solution becomes much easier. We need only find the point in the intersection of that hyper-plane with the core that maximizes the minimum difference between payment and true value. This corresponds to using the “equal-pay rule” as a secondary objective and can be implemented in practice by replacing each constraint of the form $\pi_j - m \leq \pi_j^{VCG}$ with $\pi_j + m \leq b_j(S_j)$ in the formulation $EBPO^t$, and changing the $+\varepsilon$ in the objective function to $-\varepsilon$. But without making use of the VCG payments, this alternative will not be guaranteed to minimize the maximum incentive to deviate among all other total payment minimizing points, and we therefore continue to slightly prefer the secondary objective as formulated in $EBPO^t$.

4.3 Accelerated Proxy Methods

Having established with the example of Figure 1 that the outcome of the Core Constraint Generation procedure may deviate from the threshold outcome as formulated in Hoffman et al. (2005) and Parkes (2002), and from a related Nash bargaining game solution, we now demonstrate that this technique may provide computational advantages over existing iterative proxy methods. Consider the following example from Hoffman et al. (2005):

| | | | | |
|----------------|----|-----|----|----|
| Bidder | 1 | 2 | 3 | 4 |
| Package | AB | BC* | AC | A* |
| Value | 20 | 26 | 24 | 16 |

where * in their notation denotes the winners in the efficient allocation. They present a comparison of several methods for obtaining core outcomes, and here we present their results for this problem in terms of number of rounds, along with the VCG and threshold payments, with an added row displaying the new results using Core Constraint Generation. Table 2 shows that where each of their iterative techniques ¹ may require several rounds to solve this problem, Core Constraint Generation terminates after a single round of

¹Pure Proxy refers to the algorithm described by Ausubel and Milgrom (2002), while Safe Start runs this same algorithm starting from the VCG payments. Increment Scaling uses a similar algorithm with a changing of increment, starting either from zero payments or from VCG (i.e. w/ Safe Start), proposed by Hoffman et al. (2005) who give full detail on these methods.

Table 2: Comparison of Core Constraint Generation to proxy methods

| Method | Rounds | Revenue | Payment by 2 | Payment by 4 |
|--------------------------------|--------|---------|--------------|--------------|
| Pure Proxy | 3100 | 24.02 | 12.01 | 12.01 |
| Safe Start | 800 | 24.02 | 16.01 | 8.01 |
| Increment Scaling | 20 | 24.02 | 17.01 | 7.01 |
| Increment Scaling w/Safe Start | 15 | 24.02 | 16.01 | 8.01 |
| VCG payments | - | 8.00 | 8.00 | 0.00 |
| Threshold Payments | - | 24.00 | 16.00 | 8.00 |
| Core Constraint Generation | 1 | 24.00 | 16.00 | 8.00 |

price adjustments. Starting at the VCG payments, $\pi_2^{VCG} = 8$, $\pi_4^{VCG} = 0$, we find the most violated blocking coalition consisting of just bidder 3. We then equitably divide the burden of overcoming this coalition (using EBPO^t) and find that no other blocking coalition exists. Notice that this procedure obviates the need to consider a constraint for the coalition $\{1\}$ which is made redundant by the constraint of coalition $\{3\}$.

Comparisons to all other examples worked out by Hoffman et al. (2005) verify this apparent dominance of the Core Constraint Generation technique. For every problem instance presented fully there, Core Constraint Generation terminates after a single price adjustment. We also note that in no case presented in Hoffman et al. (2005) do our computed payments differ from the threshold payments, and conclude that the phenomenon present in the example of Figure 1 is missing from the examples generated by Hoffman et al. (2005).

4.4 The Inflection Point Method

A different technique for direct computation of proxy outcomes (i.e. bidder-Pareto-optimal core outcomes) is provided by Wurman et al. (2004), who observe the existence of change points or inflection points in the price trajectories followed in the typical proxy auction as introduced by Ausubel and Milgrom (2002). At each inflection point the behavior of the ascending proxy auction changes; at a particular payment vector a bidder will no longer find it profitable to compete for the same bundle or bundles he has been pursuing for the previous rounds and will change his attention to compete for a different bundle. Wurman et al. (2004) provide a mixed-integer linear program (MIP) which traces the payment trajectory, telling at each stage which bundles a bidder is continuing to pursue between inflection points, and the allocation(s) supporting each payment vector along the trajectory. Although this research remains interesting for its ability to “jump” to the points of interest in the ascending proxy auction, we argue that for large-scale practical applications that the trajectory information is irrelevant to the final outcome. Since this additional information (illustrating how the final payments are arrived at through an ascending procedure) is achieved only through additional computational complexity, we argue that the Core Constraint Generation procedure presented here provides a more viable computational solution for applications in which intermediate payment adjustments can be

ignored.

In order to demonstrate a direct comparison of our algorithm to the inflection point method, we apply our technique on the following worked example of Wurman et al. (2004):

| | A | B | AB | C | AC | BC | ABC |
|---------|----|---|----|---|----|----|-----|
| Buyer 1 | 10 | 3 | 18 | 2 | 18 | 10 | 20 |
| Buyer 2 | 4 | 9 | 15 | 3 | 12 | 18 | 20 |
| Buyer 3 | 1 | 3 | 11 | 9 | 16 | 17 | 25 |
| Buyer 4 | 7 | 7 | 16 | 7 | 16 | 16 | 20 |

This auction has two efficient solutions (both recognized by the inflection point method) with the following outcomes (with payments computed by the inflection point method): Buyers 1, 2, and 3 get A, B, and C, respectively, with $\pi_1 = \pi_2 = 8$, $\pi_3 = 9$; and Buyer 1 gets A, Buyer 2 gets B and C with $\pi_1 = 8$, $\pi_2 = 17$. The process solves 11 MIPs in order to achieve these solutions, each one more computationally complex than the winner determination problem for the auction.

The Core Constraint Generation algorithm first selects a particular efficient solution, so for the sake of comparison we ran our algorithm twice, once for each efficient solution. For the first (with three winners) we solve the initial winner determination IP, followed by the solution of three winner determination problems (one with each of the three winners removed) in order to determine the VCG payments. We then solve a single instance of the separation problem SEP^t , adjust prices once from the VCG payments, and terminate at the solution $\pi_1 = 7.5$, $\pi_2 = 8.5$, $\pi_3 = 9$. In total 5 IPs must be solved, each computationally equivalent to the winner determination problem. For the alternate efficient outcome, we again find the final payments after just a single payment adjustment. For this instance we solve just 4 IPs to arrive at the final payments, $\pi_1 = 7.5$, $\pi_2 = 17.5$, with one less IP because there is one less VCG payment to compute.

We first observe that although the total number of IPs/MIPs to find *all* solutions is similar for this example (9 vs. 11), in practice only a single auction outcome is necessary, so that the comparison becomes 4 or 5 vs. 11. In addition, each instance of the inflection point MIP effectively finds every solution to the winner determination problem for a particular vector of information released to the proxy, making it *necessarily more complex* than each IP solved by our algorithm, which finds a single solution to the winner determination problem.

Secondly, we observe that the Core Constraint Generation algorithm determines additional useful information, and hence achieves a different (and better) set of final prices. In particular, VCG payments are computed explicitly, allowing us to employ Theorem 2.1 and more effectively deter deviation from the honest

revelation strategy. Although both techniques minimize total payments within the core for this example, consequently minimizing the total incentive to deviate via Corollary 2.2, the secondary objective term in $EBPO^t$ allows us to minimize the maximum incentive to deviate for a single player, as well. Comparing the outcomes of the two techniques for the first efficient solution, with $\pi_1^{VCG} = 7$, $\pi_2^{VCG} = 8$, $\pi_3^{VCG} = 9$, the payment rule instituted under the inflection point method provides Buyer 1 with the opportunity to benefit up to $8 - 7 = 1$ monetary unit by deviating from truthful revelation, while the Core Constraint Generation technique divides the incentive, giving Buyers 1 and 2 each the ability to gain at most $7.5 - 7 = 8.5 - 8 = 0.5$ monetary units through unilateral deviation from honesty. Similarly for the other efficient solution, $\pi_1^{VCG} = 7$, $\pi_2^{VCG} = 17$, Buyer 1 has the incentive to deviate bounded by 1 under the inflection point method, while the Core Constraint Generation technique gives an incentive to deviate of at most 0.5 to each winning bidder.

This example illustrates that computing the VCG payments allows us to get a firm handle on the incentive properties of our auction mechanism. By spending the computational energy to solve a few winner determination problems at the outset, we are able to spread the incentive to deviate amongst the bidders so that each individual's decision to deviate is less attractive, while simultaneously making sure that the total incentive to deviate is minimized. Further, starting the Core Constraint Generation algorithm at the VCG payments allows us to guarantee that VCG payments are chosen as the final outcome whenever they are in the core. As the example of Wurman et al. (2004) shows, the ascending proxy and inflection point methods do not necessarily compute the VCG payments, and may therefore produce a solution with an inferior distribution of deviation incentives.

Using the VCG payments (by explicitly computing them) in the ascending proxy auction results in the techniques presented by Hoffman et al. (2005), which did not perform nearly as well as the Core Constraint Generation for any instance investigated. More importantly though, the added complexity of the MIP introduced by Wurman et al. (2004) is overkill; every optimal solution to an \mathcal{NP} -hard problem winner determination problem is found for several payment vectors. Further, many of the iterations of the procedure reveal information that is irrelevant to the final auction outcome. For example, before the first inflection point in an ascending proxy implementation (starting from zero payments) of the previous example, the bundle price of the set $\{A,B,C\}$ rises until it reaches 2, at which point Buyer 1 and Buyer 2 become equally interested in other bundles containing just two items. Clearly, however, this information is irrelevant to the auction outcome; if Buyer 1 and 2's values for the set $\{A,B,C\}$ are both reduced to 18, the auction outcome does not change. Solving a complex mixed-integer programming problem to determine where bidders will give up on the grand combination seems to be wasted computational effort.

The strength of the ascending proxy technique (with the accelerating techniques of Hoffman et al. (2005),

the inflection point method of Wurman et al. (2004), or otherwise) is that an auction with iterative revelation of demand is either performed explicitly or simulated. But, since the information is submitted to a proxy (and cannot be changed by the bidder) iterative revelation of demand is strategically irrelevant, and merely makes a nice story to explain the process to bidders. We propose the Core Constraint Generation procedure as an alternative that bypasses the iterative revelation of demand, in order to achieve the final outcome more directly. Only allocations that affect the final payments of the winning bidders are considered; many allocations that are only feasible at intermediate, irrelevant payment vectors can be ignored.

4.5 Variations on the Core Constraint Generation Algorithm

At no point in the development of the Core Constraint Generation algorithm was it necessary for the separation procedure to begin at the VCG payment vector. In truth, the separation problem SEP^t will find the most violated coalition of bidders starting with an efficient allocation and *any* payment vector. This provides the possibility of many variations of the algorithm, each starting from a different set of payments. Motivated by the approach of Hoffman et al. (2005), and by the beneficial incentive properties guaranteed by Theorem 2.1 and Corollary 2.2, we propose that the VCG payments provide the strongest starting point for the separation portion of the algorithm, but are open to the possibility of an alternative starting point.

Is it possible that by starting at a different set of initial payments that we reduce the overall computational burden as measured by the number of IPs solved? If we save computational time by not computing the VCG payments will total computational burden be reduced? The following example demonstrates that the answer may be ambiguous, even for a single auction.

The most natural alternative starting point to consider is the zero payment vector, and we tested this variation of the Core Constraint Generation algorithm on the inflection point example from Wurman et al. (2004). Since the VCG payments are not available, we must alter the formulation $EBPO^t$, giving it some other basis for selecting among all possible payment vectors which minimize the sum of payments within the core. Though there are several possibilities, we used the most simple: simply replace the VCG payments in the $EBPO^t$ formulation with zeros, so that the algorithm finds a payment vector that minimizes the maximum payment of any bidder among all payment vectors that minimize total payments.

We ran this alternative algorithm twice on the example from Wurman et al. (2004), again, once for each efficient solution. For the first efficient solution (with three winners), we solve 6 IPs: the initial winner determination problem, and 5 instances of the separation problem. Compared with the 5 total IPs solved when using the VCG starting point, this technique seems inferior. Surprisingly however, if we choose the alternative efficient solution for the same problem (with just two winners), we solve only 3 IPs, compared

with 4 when starting at the VCG payments. Clearly there is no dominance for one technique over the other in terms of computational burden alone, even for the same auction! (The results are identical in both final payments and number of IPs solved if we use the equal-pay secondary objective discussed in §4.2.)

Interestingly though, starting at payments of zero and minimizing the maximum difference from zero payment over all bidders, we find exactly the same payments as we do using the inflection point method (and hence the generic ascending proxy auction), for this example. As noted earlier, this solution has inferior incentive properties, causing us again to lean towards the use of a VCG starting point. The less costly starting point doesn't always lead to overall improvement of performance and forfeits the ability to approximate VCG payments explicitly.

5 Concluding Remarks

Recent trends indicate that combinatorial proxy auctions, in which bidders submit bidding information on many bundles to a proxy bidding on their behalf, are fast becoming viable for several large scale business-to-business and governmental auctions. Because of their adherence to efficient allocations, core outcomes which cannot be challenged by disgruntled bidders, and bidder-Pareto-optimality, guaranteeing that bidders do not suffer high payments unnecessarily, proxy auctions are especially attractive solutions for the allocation of public goods. Ausubel et al. (2005) provide the practical clock-proxy-auction design, motivated by a great deal of real experience conducting high-stakes combinatorial auctions. Further, this auction format has received a great deal of attention as a candidate for existing and proposed governmental auction applications (Kwerel, 2004; Cramton, 2005). Consequently, the improvements described in this paper addressing critical issues related to the proxy auction promise to have a wide-ranging impact.

The clock-proxy auction incorporates a demand revelation phase which is followed by (and terminates with) a proxy auction, which is essentially a sealed-bid combinatorial auction with a particular payment rule. This proxy phase can be run a number of different ways, but as laid out in the earlier ascending proxy auction of Ausubel and Milgrom (2002), there are only a few essential features: efficiency based on reported values, and the use of a payment rule that finds a bidder-Pareto-optimal (as opposed to a VCG or pay-as-bid) point in the core. In this paper we have presented a new “direct” method of finding auction outcomes which meet these same criteria, promising to deliver faster results in the proxy environment, to allow for auctions with a greater number of items (relevant, for example, to the large number of landing-slots in proposed FAA applications), and to assure bidders of a transparent paradigm for fair payments.

Along the way, we refined several key concepts governing these auction outcomes. A few general results indicated the virtue of the approximate VCG mechanism, justifying an approach which minimizes total

payments within the core primarily, and minimizes the maximum difference from VCG payments among all total payment minimizing solutions as a secondary objective. Our treatment provides a clear understanding of which bidder-Pareto-optimal point should be chosen from the core, where in the past this notion was somewhat vague.

The example of Figure 1 (with Proposition 4.2 and Corollary 2.2) justifies our proposal that total payment minimization within the core should be the primary objective of payment determination with payment equity following as a secondary objective, used to distinguish between multiple optima. Reworking this example with the techniques of Parkes (2002) and Nash (1950), we see that there is indeed a discrepancy, with these other techniques failing to minimize total payments, and thus not minimizing total opportunity to gain from deviation, via Corollary 2.2. Adding further support for a technique that minimizes total payments within the core, Propositions 4.1 and 4.2 describe a form of collusion that is nullified using total payment minimization, and we demonstrated how bidders can collude to achieve a “socially unacceptable” outcome using the alternative techniques.

As a final remark on the selection of a core outcome, we note that the public perception of fairness may be another strong motivation to adopt a total-payment-minimization mechanism. For governmental auctions in particular, many will take comfort that the auctioneer is not selecting an outcome according to self interest, but in the combined best interests of the bidders. Indeed, the Air Transportation Association, representing the concerns of airlines who would bid in the recently proposed auctions for airport landing-slots, voice a concern that “slot auctions should not be designed to maximize payments from airlines (Airline Business Report, 2005).” What could be a more accommodating response to these public demands than a mechanism assuring total-payment minimization over the space of socially acceptable outcomes? In private sector B2B auctions, on the other hand, assurance of an auction that reaps benefits in a minimal fashion may be the best method to maintain long-term trade relationships and encourage repeated participation. In both cases, only the core constraints, representing competition, should drive up the payments, not the artificial device used to determine payments. Looking back at the classical auction literature, we see that the second-price single-item auction also minimizes payments with respect to core constraints, reinforcing our payment mechanism as a viable generalization of the second-price auction for multiple goods.

Additionally, the Core Constraint Generation method offers computational benefits over other techniques described in the recent literature. By solving the separation problem at each point in payment space, we only consider coalitions that *actually* threaten to block a potential outcome, obviating the need to apply the entire exponential set of constraints. Our definition of core constraints (as in CORE2) are equivalent to those defined as core constraints in Hoffman et al. (2005) and Parkes (2002) through a linear change of variables, but reformulating the constraints in terms of payments, rather than payoffs, provides the proper intuition for

our observations on coalitional contribution, facilitating the formulation of the separation problem. To our knowledge, the formulation of the separation problem for violated core constraints and the development of a price adjustment procedure utilizing this approach are novel contributions to the literature. This separation paradigm has proven useful elsewhere for the solution of optimization problems with an exponential number of constraints (e.g. traveling salesman problems), and we hope that future experiments will continue to verify this for Core Constraint Generation.

Further, the Core Constraint Generation algorithm takes place without the need for a price increment as in the ascending proxy technique, accelerated or otherwise. EBPO^t upholds every generated core constraint for the remainder of the procedure. This ensures that the same coalition will not appear repeatedly as blocking in our procedure. Experience with the ascending proxy method, on the other hand, shows that after an increment and re-resolution of the winner determination problem, the same coalition may appear for several iterations.

Wurman et al. (2004) find a different way around this issue with the use of a MIP that jumps to each “change point” or “inflection point” at which the relevant coalitions change, but we argued against that technique for several reasons. By reworking the example presented by Wurman et al. (2004), we showed that the outcome arrived at by the inflection point method (and hence the Ausubel and Milgrom, 2002, method) may have final payments with an inferior distribution of incentives to deviate. Interestingly, a variation of our technique produces the same outcome, showing the benefits of “pre-processing” to find VCG payments. More importantly, the inflection point method finds more information than is needed for most practical implementations, at the cost of additional computational complexity.

Our core constraint generation algorithm shows much promise as the method of choice for the determination of bidder-Pareto-optimal core outcomes in public sector combinatorial auctions. As the examples provided in this paper show, our technique clearly dominates the proxy-auction algorithm, the accelerated proxy developed by Hoffman et al. (2005), and the inflection point approach of Wurman et al. (2004). We hope that ongoing studies (performed by the authors and by researchers associated with the FAA and FCC) testing large scale implementations will confirm the rapidness, robustness, and scalability of our approach; thereby providing a practical methodology for computing fair solutions in public sector combinatorial auctions.

References

Airline Business Report 2005. Gulf between legacy carriers, auction proponents revealed by mock proposal. Vol. 23 No. 2.

- Ausubel, L. 2000. An efficient dynamic auction for heterogeneous commodities. Working paper, University of Maryland.
- Ausubel, L., P. Cramton, and P. Milgrom 2005. The clock-proxy auction: A practical combinatorial auction design. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 5. MIT Press, forthcoming.
- Ausubel, L. and P. Milgrom 2002. Ascending auctions with package bidding. *Frontiers of Theoretical Economics* **1**(1).
- Bazaraa, M., H. Sherali, and C. Shetty 1979. *Nonlinear Programming: Theory and Algorithms*. Wiley and Sons.
- Clarke, E. 1989. Multipart pricing of public goods. *Public Choice* **8**, 19–33.
- Cramton, P. 2005. The clock-proxy auction: A practical combinatorial auction design. Oral Presentation. Dagstuhl Conference: Computing and Markets.
- de Vries, S., J. Schummer, and R. Vohra 2003. On ascending Vickrey auctions for heterogeneous objects. Working paper, Northwestern University.
- Groves, T. 1973. Incentives in teams. *Econometrica* **41**, 617–631.
- Günlük, O., L. Ladányi, and S. de Vries 2005. A branch-and-price algorithm and new test problems for spectrum auctions. *Management Science* March 2005.
- Hoffman, K., D. Menon, S. van den Heever, and T. Wilson 2005. Observations and near-direct implementations of the ascending proxy auction. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 17. MIT Press, forthcoming.
- Kwerel, E. 2004. Personal communication. NEXTOR Workshop: Government, the Airline Industry and the Flying Public: A New Way of Doing Business.
- Nash, J. 1950. The bargaining problem. *Econometrica* **18**(2), 155–162.
- Parkes, D. 2002. Notes on indirect and direct implementations of core outcomes in combinatorial auctions. Technical report, Harvard University.
- Parkes, D., J. Kalagnanam, and M. Eso 2001. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *IJCAI'01*, pages 1161–1168.

- Rothkopf, M. and R. Harstad 1995. Two models of bid-taker cheating in Vickrey auctions. *Journal of Business* **68**, 257–267.
- Rothkopf, M., A. Pekeč, and R. Harstad 1998. Computationally manageable combinatorial auctions. *Management Science* **44**, 1131–1147.
- Rothkopf, M., T. Teisberg, and E. Kahn 1990. Why are Vickrey auctions rare? *Journal of Political Economy* **98**, 94–109.
- Sakurai, Y., M. Yokoo, and S. Matsubara 1999. An efficient approximate algorithm for winner determination in combinatorial auctions. In *ACM Conference on Electronic Markets*, pages 30–37. (EC-00). ACM Press.
- Vickrey, W. 1961. Counterspeculation, auctions, and competitive sealed tenders. *J. of Finance* **16**, 8–37.
- Wurman, P., J. Zhong, and G. Cai 2004. Computing price trajectories in combinatorial auctions with proxy bidding. Working paper, Computer Science Department, North Carolina State University.