Robust Solutions for Combinatorial Auctions *

Alan Holland
Cork Constraint Computation Centre
Department of Computer Science
University College Cork, Ireland
a.holland@4c.ucc.ie

Barry O'Sullivan
Cork Constraint Computation Centre
Department of Computer Science
University College Cork, Ireland
b.osullivan@4c.ucc.ie

ABSTRACT

Bids submitted in auctions are usually treated as enforceable commitments in most bidding and auction theory literature. In reality bidders often withdraw winning bids before the transaction when it is in their best interests to do so. Given a bid withdrawal in a combinatorial auction, finding an alternative repair solution of adequate revenue without causing undue disturbance to the remaining winning bids in the original solution may be difficult or even impossible. We have called this the "*Bid-taker's Exposure Problem*". When faced with such unreliable bidders, it is preferable for the bid-taker to preempt such uncertainty by having a solution that is *robust* to bid withdrawal and provides a guarantee that possible withdrawals may be repaired easily with a bounded loss in revenue.

In this paper, we propose an approach to addressing the Bid-taker's Exposure Problem. Firstly, we use the Weighted Super Solutions framework [13], from the field of constraint programming, to solve the problem of finding a robust solution. A weighted super solution guarantees that any subset of bids likely to be withdrawn can be repaired to form a new solution of at least a given revenue by making limited changes. Secondly, we introduce an auction model that uses a form of leveled commitment contract [26, 27], which we have called *mutual bid bonds*, to improve solution reparability by facilitating backtracking on winning bids by the bid-taker. We then examine the trade-off between robustness and revenue in different economically motivated auction scenarios for different constraints on the revenue of repair solutions. We also demonstrate experimentally that fewer winning bids partake in robust solutions, thereby reducing any associated overhead in dealing with extra bidders. Robust solutions can also provide a means of selectively discriminating against distrusted bidders in a measured manner.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; K.4.4 [**Computers and Society**]: Electronic Commerce; I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, and Search.

General Terms

Algorithms, Economics, Reliability.

Keywords

Combinatorial Auctions, Bid withdrawal, Robustness, Constraint Programming, Weighted Super Solutions.

1. INTRODUCTION

A combinatorial auction (CA) [5] provides an efficient means of allocating multiple distinguishable items amongst bidders whose perceived valuations for combinations of items differ. Such auctions are gaining in popularity and there is a proliferation in their usage across various industries such as telecoms, B2B procurement and transportation [11, 19].

Revenue is the most obvious optimization criterion for such auctions, but another desirable attribute is *solution robustness*. In terms of combinatorial auctions, a robust solution is one that can withstand bid withdrawal (a *break*) by making changes easily to form a repair solution of adequate revenue. A *brittle solution* to a CA is one in which an unacceptable loss in revenue is unavoidable if a winning bid is withdrawn. In such situations the bid-taker may be left with a set of items deemed to be of low value by all other bidders. These bidders may associate a higher value for these items if they were combined with items already awarded to others, hence the bid-taker is left in an undesirable local optimum in which a form of backtracking is required to reallocate the items in a manner that results in sufficient revenue. We have called this the "Bid-taker's Exposure Problem" that bears similarities to the "Exposure Problem" faced by bidders seeking multiple items in separate single-unit auctions but holding little or no value for a subset of those items.

However, reallocating items may be regarded as disruptive to a solution in many real-life scenarios. Consider a scenario where procurement for a business is conducted using a CA. It would be highly undesirable to retract contracts from a group of suppliers because of the failure of a third party. A robust solution that is tolerant of such breaks is preferable. Robustness may be regarded as a preventative measure protecting against future uncertainty by sacrificing revenue in place of solution stability and reparability. We assume a probabilistic approach whereby the bid-taker has knowledge of the reliability of bidders from which the likelihood of an incomplete transaction may be inferred.

Repair solutions are required for bids that are seen as *brittle* (*i.e.* likely to break). Repairs may also be required for *sets* of bids deemed brittle. We propose the use of the Weighted Super

*This work has received support from Science Foundation Ireland under grant number 00/PI.1/C075. The authors wish to thank Brahim Hnich and the anonymous reviewers for their helpful comments.

EC'05, June 5–8, 2005, Vancouver, British Columbia, Canada.

Solutions (WSS) framework [13] for constraint programming, that is ideal for establishing such robust solutions. As we shall see, this framework can enforce constraints on solutions so that possible breakages are reparable.

This paper is organized as follows. Section 2 presents the Winner Determination Problem (WDP) for combinatorial auctions, outlines some possible reasons for bid withdrawal and shows how simply maximizing expected revenue can lead to intolerable revenue losses for risk-averse bid-takers. This motivates the use of robust solutions and Section 3 introduces a constraint programming (CP) framework, Weighted Super Solutions [13], that finds such solutions. We then propose an auction model in Section 4 that enhances reparability by introducing mandatory *mutual bid bonds*, that may be seen as a form of leveled commitment contract [26, 27]. Section 5 presents an extensive empirical evaluation of the approach presented in this paper, in the context of a number of well-known combinatorial auction distributions, with very encouraging results. Section 6 discusses possible extensions and questions raised by our research that deserve future work. Finally, in Section 7 a number of concluding remarks are made.

2. COMBINATORIAL AUCTIONS

Before presenting the technical details of our solution to the "Bid-taker's Exposure Problem", we shall present a brief survey of combinatorial auctions and existing techniques for handling bid withdrawal.

Combinatorial auctions involve a single bid-taker allocating multiple distinguishable items amongst a group of bidders. The bid-taker has a set of m items for sale, $M = \{1, 2, \ldots, m\}$, and bidders submit a set of bids, $\mathcal{B} = \{B_1, B_2, \ldots, B_n\}$. A bid is a tuple $B_j = \langle S_j, p_j \rangle$ where $S_j \subseteq M$ is a subset of the items for sale and $p_j \geq 0$ is a price. The WDP for a CA is to label all bids as either winning or losing so as to maximize the revenue from winning bids without allocating any item to more than one bid. The following is the integer programming formulation for the WDP:

$$max \sum_{j=1}^{n} p_j x_j \;\; s.t. \sum_{j|i \in S_j} x_j \leq 1, \; \forall i \in \{1 \ldots m\}, \;\; x_j \in \{0, 1\}.$$

This problem is $\mathcal{NP}$-complete [23] and inapproximable [25], and is otherwise known as the Set Packing Problem. The above problem formulation assumes the notion of *free disposal*. This means that the optimal solution need not necessarily sell all of the items. If the auction rules stipulate that all items *must* be sold, the problem becomes a Set Partition Problem [5]. The WDP has been extensively studied in recent years. The fastest search algorithms that find optimal solutions (*e.g.* CABOB [25]) can, in practice, solve very large problems involving thousands of bids very quickly.

2.1 The Problem of Bid Withdrawal

We assume an auction protocol with a three stage process involving the submission of bids, winner determination, and finally a transaction phase. We are interested in bid withdrawals that occur between the announcement of winning bids and the end of the transaction phase. All bids are valid until the transaction is complete, so we anticipate an expedient transaction process[1].

[1]In some instances the transaction period may be so lengthy that consideration of non-winning bids as still being valid may not be fair. Breaks that occur during a lengthy transaction phase are more difficult to remedy and may require a subsequent auction. For example, if the item is a service contract for a given period of time and the break occurs after partial fulfilment of this contract, the other bidders' valuations for the item may have decreased in a non-linear fashion.

An example of a winning bid withdrawal occurred in an FCC spectrum auction [32]. Withdrawals, or *breaks*, may occur for various reasons. Bid withdrawal may be instigated by the bid-taker when Quality of Service agreements are broken or payment deadlines are not met. We refer to bid withdrawal by the bid-taker as item withdrawal in this paper to distinguish between the actions of a bidder and the bid-taker. Harstad and Rothkopf [8] outlined several possibilities for breaks in single item auctions that include:

1. an erroneous initial valuation/bid;
2. unexpected events outside the winning bidder's control;
3. a desire to have the second-best bid honored;
4. information obtained or events that occurred after the auction but before the transaction that reduces the value of an item;
5. the revelation of competing bidders' valuations infers reduced profitability, a problem known as the "*Winner's Curse*".

Kastner *et al.* [15] examined how to handle perturbations given a solution whilst minimizing necessary changes to that solution. These perturbations may include bid withdrawals, change of valuation/items of a bid or the submission of a new bid. They looked at the problem of finding incremental solutions to restructure a supply chain whose formation is determined using combinatorial auctions [30]. Following a perturbation in the optimal solution they proceed to impose involuntary item withdrawals from winning bidders. They formulated an incremental integer linear program (ILP) that sought to maximize the valuation of the repair solution whilst preserving the previous solution as much as possible.

2.2 Being Proactive against Bid Withdrawal

When a bid is withdrawn there may be constraints on how the solution can be repaired. If the bid-taker was freely able to revoke the awarding of items to other bidders then the solution could be repaired easily by reassigning all the items to the optimal solution without the withdrawn bid. Alternatively, the bidder who reneged upon a bid may have all his other bids disqualified and the items could be reassigned based on the optimum solution without that bidder present. However, the bid-taker is often unable to freely reassign the items already awarded to other bidders. When items cannot be withdrawn from winning bidders, following the failure of another bidder to honor his bid, repair solutions are restricted to the set of bids whose items only include those in the bid(s) that were reneged upon. We are free to award items to any of the previously unsuccessful bids when finding a repair solution.

When faced with uncertainty over the reliability of bidders a possible approach is to *maximize expected revenue*. This approach does not make allowances for risk-averse bid-takers who may view a small possibility of very low revenue as unacceptable.

Consider the example in Table 1, and the optimal expected revenue in the situation where a single bid may be withdrawn. There are three submitted bids for items A and B, the third being a combination bid for the pair of items at a value of 190. The optimal solution has a value of 200, with the first and second bids as winners. When we consider the probabilities of failure, in the fourth column, the problem of which solution to choose becomes more difficult.

Computing the expected revenue for the solution with the first and second bids winning the items, denoted $\langle 1, 1, 0 \rangle$, gives:

$$(200 \times 0.9 \times 0.9) + (2 \times 100 \times 0.9 \times 0.1) + (190 \times 0.1 \times 0.1) = 181.90.$$

Table 1: Example Combinatorial Auction.

Bids	A	B	AB	Withdrawal prob
	Items			
x_1	100	0	0	0.1
x_2	0	100	0	0.1
x_3	0	0	190	0.1

If a single bid is withdrawn there is probability of 0.18 of a revenue of 100, given the fact that we cannot withdraw an item from the other winning bidder. The expected revenue for $\langle 0, 0, 1\rangle$ is:

$$(190 \times 0.9) + (200 \times 0.1) = 191.00.$$

We can therefore surmise that the second solution is preferable to the first based on expected revenue.

Determining the maximum expected revenue in the presence of such uncertainty becomes computationally infeasible however, as the number of brittle bids grows. A WDP needs to be solved for all possible combinations of bids that may fail. The possible loss in revenue for breaks is also not tightly bounded using this approach, therefore a large loss may be possible for a small number of breaks.

Consider the previous example where the bid amount for x_3 becomes 175. The expected revenue of $\langle 1, 1, 0\rangle$ (181.75) becomes greater than that of $\langle 0, 0, 1\rangle$ (177.50). There are some bid-takers who may prefer the latter solution because the revenue is never less than 175, but the former solution returns revenue of only 100 with probability 0.18. A risk-averse bid-taker may not tolerate such a possibility, preferring to sacrifice revenue for reduced risk.

If we modify our repair search so that a solution of *at least* a given revenue is guaranteed, the search for a repair solution becomes a satisfiability test rather than an optimization problem. The approaches described above are in contrast to that which we propose in the next section. Our approach can be seen as preventative in that we find an *initial allocation* of items to bidders which is robust to bid withdrawal. Possible losses in revenue are bounded by a fixed percentage of the true optimal allocation. Perturbations to the original solution are also limited so as to minimize disruption. We regard this as the ideal approach for real-world combinatorial auctions.

DEFINITION 1 (ROBUST SOLUTION FOR A CA). *A robust solution for a combinatorial auction is one where any subset of successful bids whose probability of withdrawal is greater than or equal to α can be repaired by reassigning items at a cost of at most β to other previously losing bids to form a repair solution.*

Constraints on acceptable revenue, *e.g.* being a minimum percentage of the optimum, are defined in the problem model and are thus satisfied by all solutions. The maximum cost of repair, β, may be a fixed value that may be thought of as a fund for compensating winning bidders whose items are withdrawn from them when creating a repair solution. Alternatively, β may be a function of the bids that were withdrawn. Section 4 will give an example of such a mechanism. In the following section we describe an ideal constraint-based framework for the establishment of such robust solutions.

3. FINDING ROBUST SOLUTIONS

In constraint programming [4] (CP), a constraint satisfaction problem (CSP) is modeled as a set of n variables $X = \{x_1, \ldots, x_n\}$, a set of domains $D = \{D(x_1), \ldots, D(x_n)\}$, where $D(x_i)$ is the set of finite possible values for variable x_i and a set $C = \{C_1, \ldots, C_m\}$ of constraints, each restricting the assignments of some subset of the variables in X. Constraint satisfaction involves finding values for each of the problem variables such that all constraints are satisfied. Its main advantages are its declarative nature and flexibility in tackling problems with arbitrary side constraints. Constraint optimization seeks to find a solution to a CSP that optimizes some objective function. A common technique for solving constraint optimization problems is to use branch-and-bound techniques that avoid exploring sub-trees that are known not to contain a better solution than the best found so far. An initial bound can be determined by finding a solution that satisfies all constraints in C or by using some heuristic methods.

A classical *super solution* (SS) is a solution to a CSP in which, if a small number of variables lose their values, repair solutions are guaranteed with only a few changes, thus providing solution robustness [9, 10]. It is a generalization of both fault tolerance in CP [31] and supermodels in propositional satisfiability (SAT) [7]. An (a,b)-super solution is one in which if at most a variables lose their values, a repair solution can be found by changing at most b other variables [10].

Super solutions for combinatorial auctions minimize the number of bids whose status needs to be changed when forming a repair solution [12]. Only a particular set of variables in the solution may be subject to change and these are said to be members of the *break-set*. For each combination of brittle assignments in the break-set, a *repair-set* is required that comprises the set of variables whose values must change to provide another solution. The *cardinality* of the repair set is used to measure the cost of repair. In reality, changing some variable assignments in a repair solution incurs a lower cost than others thereby motivating the use of a different metric for determining the legality of repair sets.

The *Weighted Super Solution* (WSS) framework [13] considers the *cost* of repair required, rather than simply the number of assignments modified, to form an alternative solution. For CAs this may be a measure of the compensation penalties paid to winning bidders to break existing agreements. Robust solutions are particularly desirable for applications where unreliability is a problem and potential breakages may incur severe penalties. Weighted super solutions offer a means of expressing which variables are easily re-assigned and those that incur a heavy cost [13]. Hebrard *et al.* [9] describe how some variables may fail (such as machines in a job-shop problem) and others may not. A WSS generalizes this approach so that there is a probability of failure associated with each assignment and sets of variables whose assignments have probabilities of failure greater than or equal to a threshold value, α, require repair solutions.

A WSS measures the cost of repairing, or reassigning, other variables using *inertia* as a metric. Inertia is a measure of a variable's aversion to change and depends on its current assignment, future assignment and the breakage variable(s).

It may be desirable to reassign items to different bidders in order to find a repair solution of satisfactory revenue. Compensation may have to be paid to bidders who lose items during the formation of a repair solution. The inertia of a bid reflects the cost of changing its state. For winning bids this may reflect the necessary compensation penalty for the bid-taker to break the agreement (if such breaches are permitted), whereas for previously losing bids this is a free operation. The total amount of compensation payable to bidders may depend upon other factors, such as the cause of the break. There is a limit to how much these overall repair costs should be, and this is given by the value β. This value may not be known in advance and

Algorithm 1: WSS(int $level$, double α, double β):Boolean

```
begin
    if level > number of variables then return true
    choose unassigned variable x
    foreach value v in the domain of x do
        assign ⟨x : v⟩
        if problem is consistent then
            foreach combination of brittle assignments, A do
                if ¬reparable(A, β) then return false;
            if WSS(level+1) then return true
    unassign x
    return false
end
```

may depend upon the break. Therefore, β may be viewed as the fund used to compensate winning bidders for the unilateral withdrawal of their bids by the bid-taker. In summary, an (α,β)-WSS allows any set of variables whose probability of breaking is greater than or equal to α be repaired with changes to the original robust solution with a cost of at most β.

The depth-first search for a WSS (see pseudo-code description in Algorithm 1) maintains arc-consistency [24] at each node of the tree. As search progresses, the reparability of each previous assignment is verified at each node by extending a partial repair solution to the same depth as the current partial solution. This may be thought of as maintaining concurrent search trees for repairs. A repair solution is provided for every possible set of break variables, A. The `WSS` algorithm attempts to extend the current partial assignment by choosing a variable and assigning it a value. Backtracking may then occur for one of two reasons: we cannot extend the assignment to satisfy the given constraints, or the current partial assignment cannot be associated with a repair solution whose cost of repair is less than β should a break occur. The procedure `reparable` searches for *partial* repair solutions using backtracking and attempts to extend the last repair found, just as in (1,*b*)-super solutions [9]; the differences being that a repair is provided for a set of breakage variables rather than a single variable and the cost of repair is considered. A summation operator is used to determine the overall cost of repair. If a fixed bound upon the size of any potential break-set can be formed, the `WSS` algorithm is $\mathcal{NP}$-complete. For a more detailed description of the `WSS` search algorithm, the reader is referred to [13], since a complete description of the algorithm is beyond the scope of this paper.

EXAMPLE 1. We shall step through the example given in Table 1 when searching for a WSS. Each bid is represented by a single variable with domain values of 0 and 1, the former representing bid-failure and the latter bid-success. The probability of failure of the variables are 0.1 when they are assigned to 1 and 0.0 otherwise. The problem is initially solved using an ILP solver such as `lp_solve` [3] or `CPLEX`, and the optimal revenue is found to be 200. A fixed percentage of this revenue can be used as a threshold value for a robust solution and its repairs. The bid-taker wishes to have a robust solution so that if a single winning bid is withdrawn, a repair solution can be formed without withdrawing items from any other winning bidder. This example may be seen as searching for a (0.1,0)-weighted super solution, β is 0 because no funds are available to compensate the withdrawal of items from winning bidders. The bid-taker is willing to compromise on revenue, but only by 5%, say, of the optimal value.

Bids 1 and 3 cannot both succeed, since they both require item A, so a constraint is added precluding the assignment in which both variables take the value 1. Similarly, bids 2 and 3 cannot both win so another constraint is added between these two variables. Therefore, in this example the set of CSP variables is $V = \{x_1, x_2, x_3\}$, whose domains are all $\{0, 1\}$. The constraints are $x_1 + x_3 \leq 1$, $x_2 + x_3 \leq 1$ and $\sum_{x_i \in V} a_i x_i \geq 190$, where a_i reflects the relevant bid-amounts for the respective bid variables. In order to find a robust solution of optimal revenue we seek to maximize the sum of these amounts, $max \sum_{x_i \in V} a_i x_i$.

When all variables are set to 0 (see Figure 1(a) branch 3), this is not a solution because the minimum revenue of 190 has not been met, so we try assigning bid3 to 1 (branch 4). This is a valid solution but this variable is brittle because there is a 10% chance that this bid may be withdrawn (see Table 1). Therefore we need to determine if a repair can be formed should it break. The search for a repair begins at the first node, see Figure 1(b). Notice that value 1 has been removed from bid3 because this search tree is simulating the withdrawal of this bid. When bid1 is set to 0 (branch 4.1), the maximum revenue solution in the remaining subtree has revenue of only 100, therefore search is discontinued at that node of the tree. Bid1 and bid2 are both assigned to 1 (branches 4.2 and 4.4) and the total cost of both these changes is still 0 because no compensation needs to be paid for bids that change from losing to winning. With bid3 now losing (branch 4.5), this gives a repair solution of 200. Hence $\langle 0, 0, 1\rangle$ is reparable and therefore a WSS. We continue our search in Figure 1(a) however, because we are seeking a robust solution of *optimal* revenue.

When bid1 is assigned to 1 (branch 6) we seek a partial repair for this variable breaking (branch 5 is not considered since it offers insufficient revenue). The repair search sets bid1 to 0 in a separate search tree, (not shown), and control is returned to the search for a WSS. Bid2 is set to 0 (branch 7), but this solution would not produce sufficient revenue so bid2 is then set to 1 (branch 8). We then attempt to extend the repair for bid1 (not shown). This fails because the repair for bid1 cannot assign bid2 to 0 because the cost of repairing such an assignment would be ∞, given that the auction rules do not permit the withdrawal of items from winning bids. A repair for bid1 breaking is therefore not possible because items have already been awarded to bid2. A repair solution with bid2 assigned to 1 does not produce sufficient revenue when bid1 is assigned to 0. The inability to withdraw items from winning bids implies that $\langle 1, 1, 0\rangle$ is an irreparable solution when the minimum tolerable revenue is greater than 100. The italicized comments and dashed line in Figure 1(a) illustrate the search path for a WSS if both of these bids were deemed reparable. $\triangle$

Section 4 introduces an alternative auction model that will allow the bid-taker to receive compensation for breakages and in turn use this payment to compensate other bidders for withdrawal of items from winning bids. This will enable the reallocation of items and permit the establishment of $\langle 1, 1, 0\rangle$ as a second WSS for this example.

4. MUTUAL BID BONDS: A BACKTRACKING MECHANISM

Some auction solutions are inherently brittle and it may be impossible to find a robust solution. If we can alter the rules of an auction so that the bid-taker can retract items from winning bidders, then the reparability of solutions to such auctions may be improved. In this section we propose an auction model that permits bid and item withdrawal by the bidders and bid-taker, respectively.

We propose a model that incorporates *mutual bid bonds* to enable solution reparability for the bid-taker, a form of insurance against

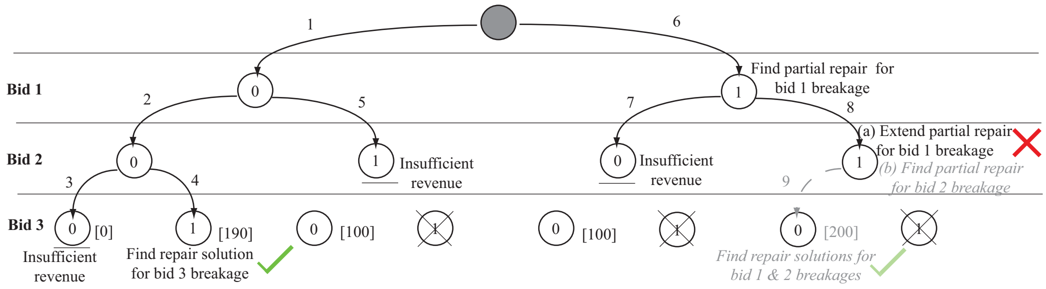

(a) Search for WSS.

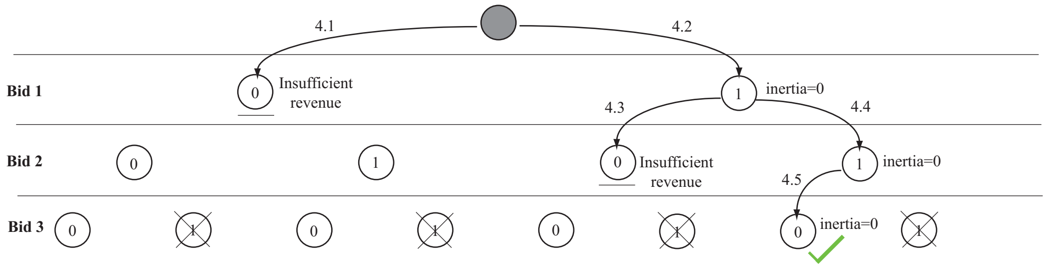

(b) Search for a repair for bid 3 breakage.

Figure 1: Search Tree for a WSS without item withdrawal.

the winner's curse for the bidder whilst also compensating bidders in the case of item withdrawal from winning bids. We propose that such "*Winner's Curse & Bid-taker's Exposure*" insurance comprise a fixed percentage, κ, of the bid amount for all bids. Such mutual bid bonds are mandatory for each bid in our model[2]. The conditions attached to the bid bonds are that the bid-taker be allowed to annul winning bids (item withdrawal) when repairing breaks elsewhere in the solution. In the interests of fairness, compensation is paid to bidders from whom items are withdrawn and is equivalent to the penalty that would have been imposed on the bidder should he have withdrawn the bid.

Combinatorial auctions impose a heavy computational burden on the bidder so it is important that the hedging of risk should be a simple and transparent operation for the bidder so as not to further increase this burden unnecessarily. We also contend that it is imperative that the bidder knows the potential penalty for withdrawal in advance of bid submission. This information is essential for bidders when determining how aggressive they should be in their bidding strategy. Bid bonds are commonplace in procurement for construction projects. Usually they are mandatory for all bids, are a fixed percentage, κ, of the bid amount and are unidirectional in that item withdrawal by the bid-taker is not permitted. Mutual bid bonds may be seen as a form of leveled commitment contract in which both parties may break the contract for the same fixed penalty. Such contracts permit unilateral decommitment for prespecified penalties. Sandholm *et al.* showed that this can increase the expected payoffs of all parties and enables deals that would be impossible under full commitment [26, 28, 29].

In practice a bid bond typically ranges between 5 and 20% of the bid amount [14, 18]. If the decommitment penalties are the same for both parties in all bids, κ does not influence the reparability of a given set of bids. It merely influences the levels of penalties and compensation transacted by agents. Low values of κ incur low bid withdrawal penalties and simulate a dictatorial bid-taker who does not adequately compensate bidders for item withdrawal. Andersson and Sandholm [1] found that myopic agents reach a higher social welfare quicker if they act selfishly rather than cooperatively when penalties in leveled commitment contracts are low. Increased levels of bid withdrawal are likely when the penalties are low also.

High values of κ tend towards full-commitment and reduce the advantages of such "Winner's Curse & Bid-taker's Exposure" insurance. The penalties paid are used to fund a reassignment of items to form a repair solution of sufficient revenue by compensating previously successful bidders for withdrawal of the items from them.

EXAMPLE 2. Consider the example given in Table 1 once more, where the bids also comprise a mutual bid bond of 5% of the bid amount. If a bid is withdrawn, the bidder forfeits this amount and the bid-taker can then compensate winning bidders whose items are withdrawn when trying to form a repair solution later. The search for repair solutions for breaks to bid1 and bid2 appear in Figures 2(a) and 2(b), respectively[3].

When bid1 breaks, there is a compensation penalty paid to the bid-taker equal to 5 that can be used to fund a reassignment of the items. We therefore set β to 5 and this becomes the maximum expenditure allowed to withdraw items from winning bidders. β may also be viewed as the size of the fund available to facilitate backtracking by the bid-taker. When we extend the partial repair for bid1 so that bid2 loses an item (branch 8.1), the overall cost of repair increases to 5, due to this item withdrawal by the bid-taker,

[2] Making the insurance optional may be beneficial in some instances. If a bidder does not agree to the insurance, it may be inferred that he may have accurately determined the valuation for the items and therefore less likely to fall victim to the winner's curse. The probability of such a bid being withdrawn may be less, so a repair solution may be deemed unnecessary for this bid. On the other hand it decreases the reparability of solutions.

[3] The actual implementation of WSS search checks previous solutions to see if they can repair breaks before searching for a new repair solution. $\langle 0, 0, 1\rangle$ is a solution that has already been found so the search for a repair in this example is not strictly necessary but is described for pedagogical reasons.

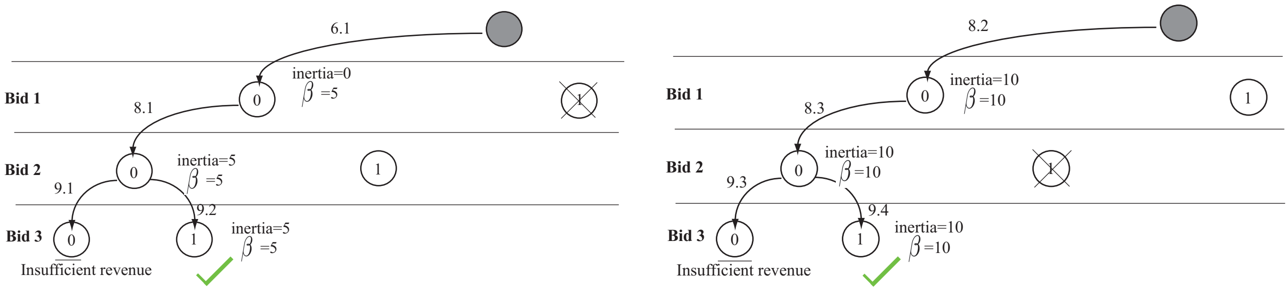

(a) Search for a repair for bid 1 breakage.

(b) Search for a repair for bid 2 breakage.

Figure 2: Repair Search Tree for breaks 1 and 2, $\kappa = 0.05$.

and is just within the limit given by β. In Figure 1(a) the search path follows the dashed line and sets bid3 to be 0 (branch 9). The repair solutions for bids 1 and 2 can be extended further by assigning bid3 to 1 (branches 9.2 and 9.4). Therefore, $\langle 1, 1, 0\rangle$ may be considered a robust solution. Recall, that previously this was not the case. $\triangle$

Using mutual bid bonds thus increases reparability and allows a robust solution of revenue 200 as opposed to 190, as was previously the case.

5. EXPERIMENTS

We have used the Combinatorial Auction Test Suite (CATS) [16] to generate sample auction data. We generated 100 instances of problems in which there are 20 items for sale and 100-2000 bids that may be dominated in some instances[4]. Such dominated bids can participate in repair solutions although they do not feature in optimal solutions. CATS uses economically motivated bidding patterns to generate auction data in various scenarios. To motivate the research presented in this paper we use sensitivity analysis to examine the brittleness of optimal solutions and hence determine the types of auctions most likely to benefit from a robust solution. We then establish robust solutions for CAs using the WSS framework.

5.1 Sensitivity Analysis for the WDP

We have performed sensitivity analysis of the following four distributions: airport take-off/landing slots (`matching`), electronic components (`arbitrary`), property/spectrum-rights (`regions`) and transportation (`paths`). These distributions were chosen because they describe a broad array of bidding patterns in different application domains.

The method used is as follows. We first of all determined the optimal solution using `lp_solve`, a mixed integer linear program solver [3]. We then simulated a single bid withdrawal and re-solved the problem with the other winning bids remaining fixed, *i.e.* there were no involuntary dropouts. The optimal repair solution was then determined. This process is repeated for all winning bids in the overall optimal solution, thus assuming that all bids are brittle. Figure 3 shows the average revenue of such repair solutions as a percentage of the optimum. Also shown is the average worst-case scenario over 100 auctions. We also implemented an auction rule that disallows bids from the reneging bidder participate in a repair[5].

Figure 3(a) illustrates how the `paths` distribution is inherently the most robust distribution since when any winning bid is withdrawn the solution can be repaired to achieve over 98.5% of the optimal revenue on average for auctions with more than 250 bids. There are some cases however when such withdrawals result in solutions whose revenue is significantly lower than optimum. Even in auctions with as many as 2000 bids there are occasions when a single bid withdrawal can result in a drop in revenue of over 5%, although the average worst-case drop in revenue is only 1%. Figure 3(b) shows how the `matching` distribution is more brittle on average than `paths` and also has an inferior worst-case revenue on average. This trend continues as the `regions-npv` (Figure 3(c)) and `arbitrary-npv` (Figure 3(d)) distributions are more brittle still. These distributions are clearly sensitive to bid withdrawal when no other winning bids in the solution may be involuntarily withdrawn by the bid-taker.

5.2 Robust Solutions using WSS

In this section we focus upon both the `arbitrary-npv` and `regions-npv` distributions because the sensitivity analysis indicated that these types of auctions produce optimal solutions that tend to be most brittle, and therefore stand to benefit most from solution robustness. We ignore the auctions with 2000 bids because the sensitivity analysis has indicated that these auctions are inherently robust with a very low average drop in revenue following a bid withdrawal. They would also be very computationally expensive, given the extra complexity of finding robust solutions.

A pure CP approach needs to be augmented with global constraints that incorporate operations research techniques to increase pruning sufficiently so that thousands of bids may be examined. Global constraints exploit special-purpose filtering algorithms to improve performance [21]. There are a number of ways to speed up the search for a weighted super solution in a CA, although this is not the main focus of our current work. Polynomial matching algorithms may be used in auctions whose bid length is short, such as those for airport landing/take-off slots for example. The integer programming formulation of the WDP stipulates that a bid either loses or wins. If we relax this constraint so that bids can partially win, this corresponds to the linear relaxation of the problem and is solvable in polynomial time. At each node of the search tree we can quickly solve the linear relaxation of the remaining problem in the subtree below the current node to establish an upper bound on remaining revenue. If this upper bound plus revenue in the parent tree is less than the current lower bound on revenue, search at that node can cease. The (continuous) LP relaxation thus provides a vital speed-up in the search for weighted super solutions, which we have exploited in our implementation. The LP formulation is as follows:

$$max \sum_{x_i \in V} a_i x_i$$

[4]The CATS flags included int_prices with the bid_alpha parameter set to 1000.

[5]We assumed that all bids in a given XOR bid with the same dummy item were from the same bidder.

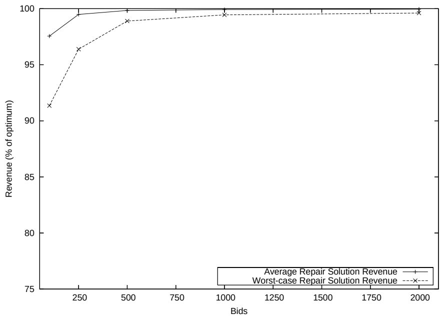

(a) paths

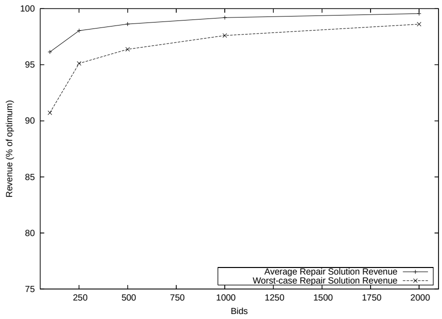

(b) matching

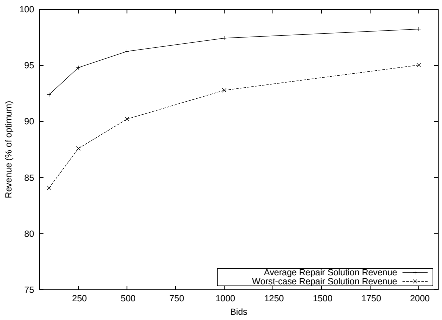

(c) regions-npv

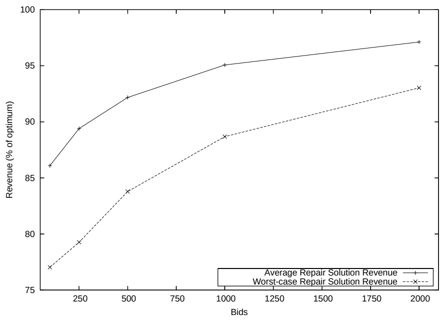

(d) arbitrary-npv

Figure 3: Sensitivity of bid distributions to single bid withdrawal.

$$s.t. \sum_{j|i \in S_j} x_j \leq 1, \ \forall i \in \{1 \ldots m\}, \ \ x_j \geq 0, \ \ x_j \in \mathbb{R}.$$

Additional techniques, that are outlined in [25], can aid the scalability of a CP approach but our main aim in these experiments is to examine the robustness of various auction distributions and consider the tradeoff between robustness and revenue. The WSS solver we have developed is an extension of the super solution solver presented in [9, 10]. This solver is, in turn, based upon the EFC constraint solver [2].

Combinatorial auctions are easily modeled as a constraint optimization problems. We have chosen the branch-on-bids formulation because in tests it worked faster than a branch-on-items formulation for the `arbitrary-npv` and `regions-npv` distributions. All variables are binary and our search mechanism uses a reverse lexicographic value ordering heuristic. This complements our dynamic variable ordering heuristic that selects the most promising unassigned variable as the next one in the search tree. We use the product of the solution of the LP relaxation and the degree of a variable to determine the likelihood of its participation in a robust solution. High values in the LP solution are a strong indication of variables most likely to form a high revenue solution whilst the a variable's degree reflects the number of other bids that overlap in terms of desired items. Bids for large numbers of items tend to be more robust, which is why we weight our robust solution search in this manner. We found this heuristic to be slightly more effective than the LP solution alone. As the number of bids in the auction increases however, there is an increase in the inherent robustness of solutions so the degree of a variable loses significance as the auction size increases.

5.3 Results

Our experiments simulate three different constraints on repair solutions. The first is that no winning bids are withdrawn by the bid-taker and a repair solution must return a revenue of at least 90% of the optimal overall solution. Secondly, we relaxed the revenue constraint to 85% of optimum. Thirdly, we allowed backtracking by the bid-taker on winning bids using mutual bid bonds but maintaining the revenue constraint at 90% of optimum.

Prior to finding a robust solution we solved the WDP optimally using `lp_solve` [3]. We then set the minimum tolerable revenue for a solution to be 90% (then 85%) of the revenue of this optimal solution. We assumed that all bids were brittle, thus a repair solution is required for every bid in the solution. Initially we assume that no backtracking was permitted on assignments of items to other winning bids given a bid withdrawal elsewhere in the solution. Table 2 shows the percentage of optimal solutions that are robust for minimum revenue constraints for repair solutions of 90% and 85% of optimal revenue. Relaxing the revenue constraint on repair solutions to 85% of the optimum revenue greatly increases the number of optimal solutions that are robust. We also conducted experiments on the same auctions in which backtracking by the bid-taker is permitted using mutual bid bonds. This significantly improves the reparability of optimal solutions whilst still maintaining repair solutions of 90% of optimum. An interesting feature of the `arbitrary-npv` distribution is that optimal solutions can become more brittle as the number of bids increases. The reason for this is that optimal solutions for larger auctions have more winning bids. Some of the optimal solutions for the smallest auctions with 100 bids have only one winning bidder. If this bid is withdrawn it is usually easy to find a new repair solution within 90% of the previous optimal revenue. Also, repair solutions for bids that contain a small number of items may be made difficult by the fact that a reduced number of bids cover only a subset of those items. A mitigating factor is that such bids form a smaller percentage of the revenue of the optimal solution on average.

We also implemented a rule stipulating that any losing bids from

Table 2: Optimal Solutions that are Inherently Robust (%).

	#Bids				
Min Revenue	100	250	500	1000	2000
`arbitrary-npv`					
repair $\geq 90\%$	21	5	3	37	93
repair $\geq 85\%$	26	15	40	87	100
MBB & repair $\geq 90\%$	41	35	60	94	≥ 93
`regions-npv`					
repair $\geq 90\%$	30	33	61	91	98
repair $\geq 85\%$	50	71	95	100	100
MBB & repair $\geq 90\%$	60	78	96	99	≥ 98

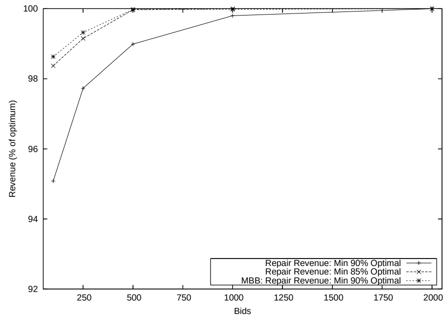

(a) regions-npv

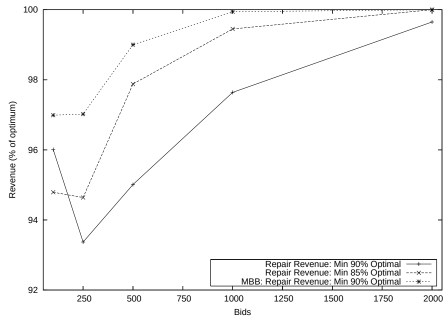

(b) arbitrary-npv

Figure 4: Revenue of optimal robust solutions.

Table 3: Occurrence of Robust Solutions (%).

	#Bids			
Min Revenue	100	250	500	1000
`arbitrary-npv`				
repair $\geq 90\%$	58	39	51	98
repair $\geq 85\%$	86	88	94	99
MBB & repair $\geq 90\%$	78	86	98	100
`regions-npv`				
repair $\geq 90\%$	61	70	97	100
repair $\geq 85\%$	89	99	99	100
MBB & repair $\geq 90\%$	83	96	100	100

a withdrawing bidder cannot participate in a repair solution. This acts as a disincentive for strategic withdrawal and was also used previously in the sensitivity analysis. In some auctions, a robust solution may not exist. Table 3 shows the percentage of auctions that support robust solutions for the `arbitrary-npv` and `regions-npv` distributions. It is clear that finding robust solutions for the former distribution is particularly difficult for auctions with 250 and 500 bids when revenue constraints are 90% of optimum. This difficulty was previously alluded to by the low percentage of optimal solutions that were robust for these auctions. Relaxing the revenue constraint helps increase the percentage of auctions in which robust solutions are achievable to 88% and 94%, respectively. This improves the reparability of all solutions thereby increasing the average revenue of the optimal robust solution. It is somewhat counter-intuitive to expect a reduction in reparability of auction solutions as the number of bids increases because there tends to be an increased number of solutions above a revenue threshold in larger auctions. The MBB auction model performs very well however, and ensures that robust solutions are achievable for such inherently brittle auctions without sacrificing over 10% of optimal revenue to achieve repair solutions.

Figure 4 shows the average revenue of the optimal robust solution as a percentage of the overall optimum. Repair solutions found for a WSS provide a lower bound on possible revenue following a bid withdrawal. Note that in some instances it is possible for a repair solution to have higher revenue than the original solution. When backtracking on winning bids by the bid-taker is disallowed, this can only happen when the repair solution includes two or more bids that were not in the original. Otherwise the repair bids would participate in the optimal robust solution in place of the bid that was withdrawn. A WSS guarantees minimum levels of revenue for repair solutions but this is not to say that repair solutions cannot be improved upon. It is possible to use an incremental algorithm to determine an optimal repair solution following a break, whilst safe in the knowledge that in advance of any possible bid withdrawal we can establish a lower bound on the revenue of a repair. Kastner *et al.* have provided such an incremental ILP formulation [15].

Mutual bid bonds facilitate backtracking by the bid-taker on already assigned items. This improves the reparability of all possible solutions thus increasing the revenue of the optimal robust solution on average. Figure 4 shows the increase in revenue of robust solutions in such instances. The revenues of repair solutions are bounded by at least 90% of the optimum in our experiments thereby allowing a direct comparison with robust solutions already found using the same revenue constraint but not providing for backtracking. It is immediately obvious that such a mechanism can significantly increase revenue whilst still maintaining solution robustness.

Table 4 shows the number of winning bids participating in optimal and optimal robust solutions given the three different constraints on repairing solutions listed at the beginning of this section. As the number of bids increases, more of the optimal overall solutions are robust. This leads to a convergence in the number of winning bids. The numbers in brackets are derived from the sensitivity analysis of optimal solutions that reveals the fact that almost all optimal solutions for auctions of 2000 bids are robust. We can therefore infer that the average number of winning bids in revenue-maximizing robust solutions converges towards that of the optimal overall solutions.

A notable side-effect of robust solutions is that fewer bids participate in the solutions. It can be clearly seen from Table 4 that when revenue constraints on repair solutions are tight, there are fewer winning bids in the optimal robust solution on average. This is particularly pronounced for smaller auctions in both distributions. This can win benefits for the bid-taker such as reduced overheads in dealing with fewer suppliers. Although MBBs aid solution repara-

Table 4: Number of winning bids.

Solution	#Bids 100	250	500	1000	2000
`arbitrary-npv`					
Optimal	3.31	5.60	7.17	9.31	10.63
Repair $\geq 90\%$	1.40	2.18	6.10	9.03	(≈ 10.63)
Repair $\geq 85\%$	1.65	3.81	6.78	9.31	(10.63)
MBB $(\geq 90\%)$	2.33	5.49	7.33	9.34	(≈ 10.63)
`regions-npv`					
Optimal	4.34	7.05	9.10	10.67	12.76
Repair $\geq 90\%$	3.03	5.76	8.67	10.63	(≈ 12.76)
Repair $\geq 85\%$	3.45	6.75	9.07	(10.67)	(12.76)
MBB $(\geq 90\%)$	3.90	6.86	9.10	10.68	(≈ 12.76)

bility, the number of bids in the solutions increases on average. This is to be expected because a greater fraction of these solutions are in fact optimal, as we saw in Table 2.

6. DISCUSSION AND FUTURE WORK

Bidding strategies can become complex in non-incentive-compatible mechanisms where winner determination is no longer necessarily optimal. The perceived reparability of a bid may influence the bid amount, with reparable bids reaching a lower equilibrium point and perceived irreparable bids being more aggressive.

Penalty payments for bid withdrawal also create an incentive for more aggressive bidding by providing a form of insurance against the winner's curse [8]. If a winning bidder's revised valuation for a set of items drops by more than the penalty for withdrawal of the bid, then it is in his best interests to forfeit the item(s) and pay the penalty. Should the auction rules state that the bid-taker will refuse to sell the items to any of the remaining bidders in the event of a withdrawal, then insurance against potential losses will stimulate more aggressive bidding. However, in our case we are seeking to repair the solution with the given bids. A side-effect of such a policy is to offset the increased aggressiveness by incentivizing reduced valuations in expectation that another bidder's successful bid is withdrawn. Harstad and Rothkopf [8] examined the conditions required to ensure an equilibrium position in which bidding was at least as aggressive as if no bid withdrawal was permitted, given this countervailing incentive to under-estimate a valuation. Three major results arose from their study of bid withdrawal in a single item auction:

1. Equilibrium bidding is more aggressive with withdrawal for sufficiently small probabilities of an award to the second highest bidder in the event of a bid withdrawal;
2. Equilibrium bidding is more aggressive with withdrawal if the number of bidders is large enough;
3. For many distributions of costs and estimates, equilibrium bidding is more aggressive with withdrawal if the variability of the estimating distribution is sufficiently large.

It is important that mutual bid bonds do not result in depressed bidding in equilibrium. An analysis of the resultant behavior of bidders must incorporate the possibility of a bidder winning an item and having it withdrawn in order for the bid-taker to formulate a repair solution after a break elsewhere. Harstad and Rothkopf have analyzed bidder aggressiveness [8] using a strictly game-theoretic model in which the only reason for bid withdrawal is the winner's curse. They assumed all bidders were risk-neutral, but surmised that it is entirely possible for the bid-taker to collect a risk premium from risk-averse bidders with the offer of such insurance. Combinatorial auctions with mutual bid bonds add an extra incentive to bid aggressively because of the possibility of being compensated for having a winning bid withdrawn by a bid-taker. This is militated against by the increased probability of not having items withdrawn in a repair solution. We leave an in-depth analysis of the sufficient conditions for more aggressive bidding for future work.

Whilst the WSS framework provides ample flexibility and expressiveness, scalability becomes a problem for larger auctions. Although solutions to larger auctions tend to be naturally more robust, some bid-takers in such auctions may require robustness. A possible extension of our work in this paper may be to examine the feasibility of reformulating integer linear programs so that the solutions are robust. Hebrard *et al.* [10] examined reformulation of CSPs for finding super solutions. Alternatively, it may be possible to use a top-down approach by looking at the k-best solutions sequentially, in terms of revenue, and performing sensitivity analysis upon each solution until a robust one is found. In procurement settings the principle of *free disposal* is often discounted and all items must be sold. This reduces the number of potential solutions and thereby reduces the reparability of each solution. The impact of such a constraint on revenue of robust solutions is also left for future work.

There is another interesting direction this work may take, namely *robust mechanism design*. Porter *et al.* introduced the notion of *fault tolerant mechanism design* in which agents have private information regarding costs for task completion, but also their probabilities of failure [20]. When the bid-taker has combinatorial valuations for task completions it may be desirable to assign the same task to multiple agents to ensure solution robustness. It is desirable to minimize such potentially redundant task assignments but not to the detriment of completed task valuations. This problem could be modeled using the WSS framework in a similar manner to that of combinatorial auctions.

In the case where no robust solutions are found, it is possible to optimize robustness, instead of revenue, by finding a solution of at least a given revenue that minimizes the probability of an irreparable break. In this manner the *least brittle* solution of adequate revenue may be chosen.

7. CONCLUSION

Fairness is often cited as a reason for choosing the optimal solution in terms of revenue only [22]. Robust solutions militate against bids deemed brittle, therefore bidders must earn a reputation for being reliable to relax the reparability constraint attached to their bids. This may be seen as being fair to long-standing business partners whose reliability is unquestioned. Internet-based auctions are often seen as unwelcome price-gouging exercises by suppliers in many sectors [6, 17]. Traditional business partnerships are being severed by increased competition amongst suppliers. Quality of Service can suffer because of the increased focus on short-term profitability to the detriment of the bid-taker in the long-term. Robust solutions can provide a means of selectively discriminating against distrusted bidders in a measured manner. As combinatorial auction deployment moves from large value auctions with a small pool of trusted bidders (*e.g.* spectrum-rights sales) towards lower value auctions with potentially unknown bidders (*e.g.* Supply Chain Management [30]), solution robustness becomes more relevant. As well as being used to ensure that the bid-taker is not left vulnerable to bid withdrawal, it may also be used to cement relationships with preferred, possibly incumbent, suppliers.

We have shown that it is possible to attain robust solutions for CAs with only a small loss in revenue. We have also illustrated how such solutions tend to have fewer winning bids than overall optimal solutions, thereby reducing any overheads associated with dealing with more bidders. We have also demonstrated that introducing mutual bid bonds, a form of leveled commitment contract, can significantly increase the revenue of optimal robust solutions by improving reparability. We contend that robust solutions using such a mechanism can allow a bid-taker to offer the possibility of bid withdrawal to bidders whilst remaining confident about post-repair revenue and also facilitating increased bidder aggressiveness.

8. REFERENCES

[1] Martin Andersson and Tuomas Sandholm. Leveled commitment contracts with myopic and strategic agents. *Journal of Economic Dynamics and Control*, 25:615–640, 2001. Special issue on Agent-Based Computational Economics.

[2] Fahiem Bacchus and George Katsirelos. EFC solver. `www.cs.toronto.edu/~gkatsi/efc/efc.html`.

[3] Michael Berkelaar, Kjell Eikland, and Peter Notebaert. lp_solve version 5.0.10.0. `http://groups.yahoo.com/group/lp_solve/`.

[4] Rina Dechter. *Constraint Processing*. Morgan Kaufmann, 2003.

[5] Sven DeVries and Rakesh Vohra. Combinatorial auctions: A survey. *INFORMS Journal on Computing*, pages 284–309, 2003.

[6] Jim Ericson. Reverse auctions: Bad idea. Line 56, Sept 2001.

[7] Matthew L. Ginsberg, Andrew J. Parkes, and Amitabha Roy. Supermodels and Robustness. In *Proceedings of AAAI-98*, pages 334–339, Madison, WI, 1998.

[8] Ronald M. Harstad and Michael H. Rothkopf. Withdrawable bids as winner's curse insurance. *Operations Research*, 43(6):982–994, November-December 1995.

[9] Emmanuel Hebrard, Brahim Hnich, and Toby Walsh. Robust solutions for constraint satisfaction and optimization. In *Proceedings of the European Conference on Artificial Intelligence*, pages 186–190, 2004.

[10] Emmanuel Hebrard, Brahim Hnich, and Toby Walsh. Super solutions in constraint programming. In *Proceedings of CP-AI-OR 2004*, pages 157–172, 2004.

[11] Gail Hohner, John Rich, Ed Ng, Grant Reid, Andrew J. Davenport, Jayant R. Kalagnanam, Ho Soo Lee, and Chae An. Combinatorial and quantity-discount procurement auctions benefit Mars Incorporated and its suppliers. *Interfaces*, 33(1):23–35, 2003.

[12] Alan Holland and Barry O'Sullivan. Super solutions for combinatorial auctions. In *Ercim-Colognet Constraints Workshop (CSCLP 04)*. Springer LNAI, Lausanne, Switzerland, 2004.

[13] Alan Holland and Barry O'Sullivan. Weighted super solutions for constraint programs, December 2004. Technical Report: No. UCC-CS-2004-12-02.

[14] Selective Insurance. Business insurance. `http://www.selectiveinsurance.com/psApps/Business/Ins/bonds.asp?bc=13.16.127`.

[15] Ryan Kastner, Christina Hsieh, Miodrag Potkonjak, and Majid Sarrafzadeh. On the sensitivity of incremental algorithms for combinatorial auctions. In *WECWIS*, pages 81–88, June 2002.

[16] Kevin Leyton-Brown, Mark Pearson, and Yoav Shoham. Towards a universal test suite for combinatorial auction algorithms. In *ACM Conference on Electronic Commerce*, pages 66–76, 2000.

[17] Associated General Contractors of America. Associated general contractors of America white paper on reverse auctions for procurement of construction. `http://www.agc.org/content/public/pdf/Member_Resources/ReverseAuctionWhitePaper.pdf`, 2003.

[18] National Society of Professional Engineers. A basic guide to surety bonds. `http://www.nspe.org/pracdiv/76-02surebond.asp`.

[19] Martin Pesendorfer and Estelle Cantillon. Combination bidding in multi-unit auctions. Harvard Business School Working Draft, 2003.

[20] Ryan Porter, Amir Ronen, Yoav Shoham, and Moshe Tennenholtz. Mechanism design with execution uncertainty. In *Proceedings of UAI-02*, pages 414–421, 2002.

[21] Jean-Charles Régin. Global constraints and filtering algorithms. In *Constraint and Integer Programming – Towards a Unified Methodology*, chapter 4, pages 89–129. Kluwer Academic Publishers, 2004.

[22] Michael H. Rothkopf and Aleksandar Pekeč. Combinatorial auction design. *Management Science*, 4(11):1485–1503, November 2003.

[23] Michael H. Rothkopf, Aleksandar Pekeč, and Ronald M. Harstad. Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131–1147, 1998.

[24] Daniel Sabin and Eugene C. Freuder. Contradicting conventional wisdom in constraint satisfaction. In A. Cohn, editor, *Proceedings of ECAI-94*, pages 125–129, 1994.

[25] Tuomas Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

[26] Tuomas Sandholm and Victor Lesser. Leveled Commitment Contracts and Strategic Breach. *Games and Economic Behavior*, 35:212–270, January 2001.

[27] Tuomas Sandholm and Victor Lesser. Leveled commitment contracting: A backtracking instrument for multiagent systems. *AI Magazine*, 23(3):89–100, 2002.

[28] Tuomas Sandholm, Sandeep Sikka, and Samphel Norden. Algorithms for optimizing leveled commitment contracts. In *Proceedings of the IJCAI-99*, pages 535–541. Morgan Kaufmann Publishers Inc., 1999.

[29] Tuomas Sandholm and Yunhong Zhou. Surplus equivalence of leveled commitment contracts. *Artificial Intelligence*, 142:239–264, 2002.

[30] William E. Walsh, Michael P. Wellman, and Fredrik Ygge. Combinatorial auctions for supply chain formation. In *ACM Conference on Electronic Commerce*, pages 260–269, 2000.

[31] Rainier Weigel and Christian Bliek. On reformulation of constraint satisfaction problems. In *Proceedings of ECAI-98*, pages 254–258, 1998.

[32] Margaret W. Wiener. Access spectrum bid withdrawal. `http://wireless.fcc.gov/auctions/33/releases/da011719.pdf`, July 2001.