

# Preference Representation for Multi-Unit Multiattribute Auctions

Yagil Engel, Kevin M. Lochner, and Michael P. Wellman

University of Michigan, Computer Science & Engineering  
2260 Hayward St, Ann Arbor, MI 48109-2121, USA  
{yagil,klochner,wellman}@umich.edu

## Abstract

The problem of multi-unit multiattribute trading, though useful in practical procurement settings, has been rarely addressed in auction literature. We present a general framework to structure specifications of preferences over multi-unit multiattribute outcomes, allowing flexible tradeoffs between expressiveness and compactness of representation. Next, we use this framework to propose an auction mechanism with useful economic and computational properties, applicable over a substantial part of the general preference domain.

## Background and introduction

Procurement auctions typically involve various non-price aspects of a deal, such as quality measures, delivery and service information, and supplier qualification criteria. A multiattribute mechanism facilitates such negotiations, and potentially achieves higher welfare by picking contract configurations that are both valued by the buyer and less costly for the supplier. Various such mechanisms have been suggested in the auction literature (Che 1993; Bichler 2001; Parkes & Kalagnanam 2005; Engel & Wellman 2007). The design and implementation of a multiattribute auction presents several technical challenges, including preference extraction, compact preference representation, and a difficult optimization problem for clearing.

Perhaps due to the inherent complexity of multiattribute auctions, nearly all prior proposals for multiattribute preference handling and mechanism design have assumed *single-sourcing*, where a single bidder is selected as supplier. As single-supplier relationships introduce a risk of cost overruns and supply disruptions or stock-outs, buyers often have a preference for *multi-sourcing*, where supply contracts are distributed among multiple winners. Limited supplier capacities may also contribute to the need for multi-sourcing. In practice, many procurement auctions have dealt with multiple units over multiattribute goods (Meti *et al.* 2005; Sandholm *et al.* 2006). To our knowledge, however, these do not support structured multiattribute preferences or provide efficiency guarantees across the space of configurations.

To address multi-sourcing issues, we consider *multi-unit multiattribute* (MUMA) auctions. The MUMA problem

combines the dimensionality of the single-unit multiattribute domain with the combinatorial complexity of multi-unit auctions of heterogeneous goods. Due to the complexity of this problem, we first restrict attention to the representation and expression of preferences over this domain. We consider the problem in its full generality, then suggest simplifications based on assumptions over preference structure.

In prior work on MUMA double auctions (Engel, Wellman, & Lochner 2006), we reasoned directly from bid expressiveness to the computational complexity of clearing, finding constraints on the bidding language which led to tractable algorithms. Our current approach takes a *preference representation* perspective, rather than bid expressiveness, which brings several advantages. It lets us capture a fully general representation of this domain (whereas we previously did not represent complementarities), it lends itself easily to the expression of the welfare maximization problem, and most importantly allows the use of well-founded tools from multiattribute utility theory to simplify the preference representation. This framework lets us use meaningful preference structure in order to achieve value decomposition, which we apply to effectively decouple the multiattribute domain from the multi-unit problem. We then leverage this result to create an auction mechanism for multi-sourcing procurement.

Some issues relevant to our work are explored in the literature on *side constraints*, which place hard constraints on the space of allocations acceptable to the bid taker in multi-object and combinatorial auctions. Work by Sandholm and Suri (2001) shows that most such constraints posed by the bid taker render the winner determination problem NP-hard. Bichler and Kalagnanam (2005) explicitly consider the problem of multi-unit multiattribute procurement auctions, focusing on the winner determination problem given various buyer-imposed constraints. Our approach is a general framework to express *preferences* over the multi-unit multiattribute domain rather than hard constraints. In that sense another relevant work is from Boutilier, Sandholm and Shields (2004), which opts for a utility-based representation over hard constraints for allocation preferences in combinatorial auctions, and focuses on the elicitation of tradeoffs between constraints and cost.

In the next section, we formally define the MUMA domain and introduce various simplifying assumptions over

preferences. We then provide computational results that improve on our previous work. Finally we present our multi-sourcing procurement auction, along with a detailed example.

## MUMA preferences

Let a good or a service be defined by a vector of attributes,  $A = (a_1, \dots, a_m)$ . These attributes may include such things as physical characteristics, supplier qualities (for procurement auctions), or contractual details (delivery date, payment terms, etc.). Let  $D_i$  denote the domain of  $a_i$ . A unique instantiation of attributes defines a *configuration*,  $\theta \in \Theta = \prod_{i=1}^m D_i$ .

Our framework is designed for the general case of two-sided markets, for which the typical one-sided procurement auction is a special case. The setting includes a set of buyers  $B$  and a set of sellers  $S$ , and the outcome of a MUMA auction is a set of bilateral trades and transfer payments. Trade  $t$  takes the form  $t = (\theta, q, b, s)$ , signifying that buyer  $b$  buys  $q > 0$  units of configuration  $\theta$  from seller  $s$ . Let  $\mathcal{T}$  denote the set of all possible trades:

$$\mathcal{T} = \{(\theta, q, b, s) \mid b \in B, s \in S, q \geq 0, \theta \in \Theta\}.$$

To simplify the presentation, the following definitions and claims are provided with respect to buyers, but apply to sellers with only minor modification.

For any global set of trades  $T \subset \mathcal{T}$ , we denote the allocation to buyer  $\hat{b}$  by  $T_{\hat{b}}$ ,

$$T_{\hat{b}} = \{(\theta, q, b, s) \in T \mid b = \hat{b}\}.$$

Note that

$$\forall T, \bigcup_{b \in B} T_b = \bigcup_{s \in S} T_s = T, \quad (1)$$

as both include all the trades in the set  $T$ .

Let  $\mathcal{T}_{\hat{b}}$  denote the set of all possible trades in which  $\hat{b}$  participates:

$$\mathcal{T}_{\hat{b}} = \{(\theta, q, b, s) \mid b = \hat{b}, (\theta, q, b, s) \in \mathcal{T}\}.$$

We define the set of all possible allocations to  $\hat{b}$ ,

$$\mathcal{O}_{\hat{b}} = 2^{\mathcal{T}_{\hat{b}}}.$$

Any element of  $\mathcal{O}_{\hat{b}}$  can be described using a non-negative vector of length  $|\Theta||S|$ , where each component designates the number of units that  $\hat{b}$  buys of a specific  $(\theta, s)$  pair (where zero quantity reflects no trade). We assume that  $\hat{b}$  has a full preference order over  $\mathcal{O}_{\hat{b}}$ .

We adopt the following common assumptions from mechanism design: free disposal, indifference to the allocation of other traders, and quasi-linearity in money of traders' preferences. Under these assumptions and standard assumptions from utility theory,  $\hat{b}$ 's preferences can be represented by a utility function  $u_{\hat{b}} : \mathcal{O}_{\hat{b}} \rightarrow \mathcal{R}$ . By quasi-linearity,  $u_{\hat{b}}$  can be scaled to monetary terms, and represent *willingness-to-pay* (wtp) per allocation. In the case of a seller  $s$ , the utility  $u_s$  is usually referred to as cost, and represented in bidding as

willingness-to-accept. We use the term wtp in referring to both buyers and sellers.<sup>1</sup>

We define the welfare optimization problem for this domain, based on buyers' and sellers' wtp.

**Definition 1.** The *global multiattribute allocation problem* (GMAP) is defined as follows:

$$GMAP = \max_{T \subseteq \mathcal{T}} \left( \sum_{b \in B} u_b(T_b) - \sum_{s \in S} u_s(T_s) \right)$$

The first restrictive assumption we make is that a trader can trade only a single configuration with each trading partner. Though non-trivial, we find this assumption reasonable for this domain. In contrast to combinatorial auctions, the central purpose of multiattribute auctions is to find surplus maximizing *configurations*, rather than surplus maximizing *combinations* of goods. The primary goal of the multi-unit generalization is to achieve an ideal supplier mix, whereas achieving a heterogeneous configuration mix is typically a secondary consideration. Heterogeneity of supply from the same supplier, if desired at all, is generally much lower in priority.

This assumption leads to an exponential reduction in dimensionality: any allocation can now be defined by a vector of length  $(|A| + 1)|S|$ , which we represent as a matrix of  $|A| + 1$  columns (one for each attribute, and one for the quantity of the trade) and  $|S|$  rows (one per trading partner). We call this matrix an *allocation matrix*, as an instance of this matrix represents an allocation. For example, in buyer  $\hat{b}$ 's matrix, row  $j$  represents a trade between  $\hat{b}$  and seller  $s_j$ , where the first  $m$  columns define the trade configuration (as values to attributes  $a_{1_j}, \dots, a_{m_j}$ ), and the attribute at the last column,  $q_j$ , defines the trade quantity. We stress that an instantiation of this matrix represents a single allocation, and that expressing wtp over all allocations requires the definition of a reserve price for every possible such instantiation.

Even with this significant simplification, the domain is still too large. The representation of utility over the allocation matrix is at best exponential in the number of elements of the matrix (assuming all attribute domains are discretized and bounded), and given a representation we cannot expect computational mechanisms to solve GMAP for more than tiny instances. However, since the domain is now representable by a restricted set of attributes, we can use tools from multiattribute utility theory to tame it further. In particular, we use the representation scheme known as *generalized additive independence* (Fishburn 1967; Bacchus & Grove 1995).

**Definition 2.** Let  $C$  be a set of attributes, and let  $I_1, \dots, I_g \subseteq C$  such that  $\bigcup_{i=1}^g I_i = C$ .  $I_1, \dots, I_g$  are called *generalized additive independent* (GAI) if there exist functions  $f_1, \dots, f_g$  such that

$$u(C) = \sum_{r=1}^g f_r(I_r). \quad (2)$$

<sup>1</sup>We use  $u_j$  when referring to a seller  $s_j$ , and drop the subscript altogether in generic statements about utility, when not referring to a particular trader.

We call each  $I_r$  a GAI element, and any assignment to  $I_r$  a sub-configuration.

In another paper (Engel & Wellman 2007), we show how a GAI structure for wtp functions can be constructed from simpler statements over preferences. In addition, for a decomposition to be useable one also needs to know the functional form of the lower dimensional functions  $f_r$ . To this end, we showed that the resulting GAI decomposition supports a functional form that is identical to the functional form of the GAI decomposition of *expected* utility functions (von Neumann-Morgenstern), as was shown by Fishburn (1967). To present this functional form we adopt the following standard notation. Let the set of attributes be  $C = (c_1, \dots, c_n)$ . Let  $(\tilde{c}_1, \dots, \tilde{c}_n)$  be a predefined vector called the *reference outcome*. For any  $I \subseteq C$ , the function  $u([I])$  designates the utility function evaluated on  $I$ , holding all attributes  $C \setminus I$  to their reference levels. Then,

$$\begin{aligned} f_1 &= u([I_1]), \text{ and for } r = 2, \dots, g & (3) \\ f_r &= u([I_r]) + \sum_{k=1}^{r-1} (-1)^k \sum_{1 \leq i_1 < \dots < i_k < r} u\left(\left[\bigcap_{s=1}^k I_{i_s} \cap I_r\right]\right) \end{aligned}$$

Though complicated on the surface, this expression bears intuitive meaning. We simply add up utilities of subsets, but to avoid double counting we subtract utilities of pairwise intersection, at which point we have over-deducted and must add back those of three-wise intersection, and so on. We explicitly apply this functional form in some of our results below.

We denote the attributes that refer to the trade with  $s_j$  by  $t_j = \{a_{1_j}, \dots, a_{m_j}, q_j\}$  (the set containing the attributes in row  $j$ ), and denote the set of attributes in the quantity column of the matrix by  $\vec{q} = \{q_1, \dots, q_{|S|}\}$ .

The functional form above relies on predefined reference values for all attributes. For our MUMA domain we define zero as the reference value for each attribute  $q_j \in \vec{q}$ . For the remaining attributes  $a_{i_j}$  we use constants  $\tilde{a}_i$ , which comprise the reference configuration  $\tilde{\theta} = (\tilde{a}_1, \dots, \tilde{a}_m)$ . This means for example that  $u_b([t_j])$  corresponds to the wtp of  $b$  for the trade  $(\theta_j, q_j, b, s_j)$ , as if it is the only trade in the allocation to  $b$ , and that  $u_b([q_j])$  is the wtp for  $q_j$  units of  $\tilde{\theta}$ , again as the only trade in the allocation. Based on this notion, we can formally introduce preferences over a single-unit allocation, which we designate by the single-unit utility function  $\mu_b : \Theta \times S \rightarrow R$ :

$$\begin{aligned} \mu_b(\hat{\theta}, s_j) &= u_b(\{a_{1_j}, \dots, a_{m_j}, q_j\}), \\ &(a_{1_j}, \dots, a_{m_j}) = \hat{\theta}, q_j = 1, \end{aligned}$$

referring to the wtp of  $b$  for an allocation that contains a single-unit trade between  $b$  and  $s_j$  over configuration  $\hat{\theta}$ . The single-unit utility of sellers is defined similarly using the notation  $\mu_s$ , or  $\mu_j$  when referring to a particular seller  $s_j$ .

**Definition 3.** A buyer  $b$  is called *configuration aggregating* if her preferences exhibit GAI structure over the following collection of subsets:

$$\{t_1, t_2, \dots, t_{|S|}, \vec{q}\}$$

In words, a configuration aggregating buyer has additive utility over different trades, subject to preferences over distribution of quantity. Though this rules out any constraints on aggregation of configurations, it does not rule out various allocation preferences as explained below. Since the collection of GAI elements in the above definition is disjoint with the exception of  $t_j \cap \vec{q} = \{q_j\}$ , when applying the functional form in Equation (3) we get a summation over all elements minus those pairwise intersections:

$$u_b(T_b) = \sum_{j=1}^{|S|} u_b([t_j]) + u_b([\vec{q}]) - \sum_{j=1}^{|S|} u_b([q_j]) \quad (4)$$

We denote the ‘‘quantity factor’’, which is the utility factor of buyer  $b$  over the GAI element  $\vec{q}$ , by  $f_b^q(\cdot)$ :

$$f_b^q(\vec{q}) = u_b([\vec{q}]) - \sum_{j=1}^{|S|} u_b([q_j]) \quad (5)$$

The interpretation of the quantity factor is that it expresses utility over the *combinations* of trades from different trading partners, whereas the factors  $u_b([t_j])$  represent utility of each trade as an allocation by itself. We later provide a full exposition of this concept.

**Definition 4.** Buyer  $b$ ’s preferences exhibit *linear trade utility* if there exist  $q^b$  and  $\bar{q}^b$  such that

$$\forall j = 1, \dots, |S|, u_b([t_j]) = g_b(q_j) \mu_b(\theta, s_j),$$

and

$$g_b(q_j) = \begin{cases} 0 & q_j < q^b \\ q_j & q^b \leq q_j \leq \bar{q}^b \\ \bar{q}^b & q_j > \bar{q}^b \end{cases}$$

$q^b$  and  $\bar{q}^b$  are called *minimal* and *maximal* trade quantity, respectively. Similarly  $q^j$  and  $\bar{q}^j$  denote the bounds for  $s_j$ .

In words, the utility of each single trade (as the only trade in the allocation) is linear in quantity, subject to strict trader-specified upper and lower bounds. In the case of sellers, the utility (cost) below the lower bound equals the cost at the lower bound, and cost above the upper bound is considered infinite. Although we resort to a hard constraint model here (in order to achieve the result of Theorem 1), we still retain the full flexibility of the quantity factor  $f_b^q(\vec{q})$ .

**Definition 5.** The *multiattribute matching problem (MMP)* is defined for traders  $b \in B$  and  $s \in S$  as follows:

$$MMP(b, s) = \max_{\theta \in \Theta} (\mu_b(\theta, s) - \mu_s(\theta, b))$$

In one-sided procurement mechanisms, the  $\theta$  chosen by *MMP* is referred to as the efficient configuration of seller  $s$ .

### MMP-GMAP decomposition

The combination of linearity in trade quantity and configuration aggregation leads to a significantly simplifying result.

**Theorem 1.** *If all traders exhibit*

1. *configuration aggregation, and*
2. *linear trade utility, with  $q_{s_j} = 0$  for each seller  $s_j$ ,*

then the solution to GMAP consists of a set of trades, each of which is a solution to MMP for its specified pair of traders. Moreover, for any preference structure that does not agree with configuration aggregation, there exists an instance of the problem for which any solution to GMAP includes a non-MMP trade.

The proof of the first claim first asserts that when  $q_{s_j} = 0$  a trade can occur only within the quantity bounds of all traders.<sup>2</sup> We can therefore assume  $g_b(q_j) = q_j$ , and then the following functional decomposition is implied by the conditions above:

$$u_b(T_b) = \sum_{j=1}^{|S|} q_j \mu_b(\theta_j, s_j) + f_b^q(\vec{q}) \quad (6)$$

And similarly for sellers. It is then easy to show that improving the bilateral single-unit solutions always improves the total surplus of GMAP.

The usefulness of this decomposition is in the fact that the global optimization problem no longer depends on the attributes, as the attributes need to be considered only in the MMP stage. We effectively separate the problem to a multiattribute single-unit problem, and a price-only multi-unit problem. This opens the way to a generalized network flow algorithm for GMAP we introduced in the past (Engel, Wellman, & Lochner 2006), and to a new procurement mechanism. The conditions of configuration aggregating and linear trade utility are fairly restrictive, however they still allow a wide variety of expressive preferences, as described next.

## Aggregation Preferences

### Aggregation over trading partners

Theorem 1 generalizes our previous result (Engel, Wellman, & Lochner 2006) by accommodating arbitrary functions over the vector of quantities, thus enabling aggregation preferences over trading partners. For example, a buyer may want to reduce the number of suppliers she deals with, but finds a hard constraint (e.g., at least  $x$  suppliers) overly restrictive, since she might want to balance her preference for more suppliers against the cost of including inefficient suppliers. The trade elements cannot express such preferences, but the quantity factor  $f_b^q(\cdot)$  can reward the total wtp based on the distribution of quantity.

Consider the form of the quantity factor in Equation (5). This factor can increase the wtp for allocations where multiple elements of  $\vec{q}$  are non-zero, which happens when the utility of all the trades taken together,  $u_b([\vec{q}])$ , is greater than the sum of the individual trade utilities,  $u_b([q_j])$ . Similarly, preferences towards a small number of suppliers results in  $u_b([\vec{q}]) < \sum_{j=1}^{|S|} u_b([q_j])$  for allocations that have too many suppliers, meaning that the quantity factor penalizes the total wtp. Note that a different choice of reference values (non-zero values for  $\vec{q}_j$ ) would not interfere with this interpretation since the non-relevant units cancel out in Equation (4).

More generally, any configuration-independent preferences over the distribution of the allocation to suppliers can

be generated by the quantity factor. These preferences cannot depend on specific configurations, because it would violate the GAI structure. For example, a buyer cannot express a preference such as “increase the number of suppliers when the quality of insurance is low”. Such a statement can be expressed using another GAI element, which includes  $\vec{q}$  and the column of the “insurance” attributes. Note that better insurance per se is rewarded through the configuration utility factors  $u_b([t_j])$ .

As mentioned, our definitions and analysis apply to sellers with minor adjustments. It is important to note however that when there is a single seller,  $f_b^q(\vec{q}) = 0$  according to (5). Similarly, in procurement mechanisms with a single buyer there are no quantity factors for sellers.

To illustrate the expressiveness of the model and its limitation we consider the following problem. A buyer  $B$  represents a procurement department that wishes to purchase LCD monitors for the company’s employees. Negotiable attributes contain contractual aspects such as warranty, delivery, insurance of deal, and payment terms. Suppliers are also evaluated on a measure of product quality and supplier reliability reputation, both accessible to  $B$ . Each offer provides product specifications over resolution, # colors, # nits, and its compatibility with wall mounting hardware.

It is reasonable to assume that  $B$  does not object to having suppliers with different contractual attributes and different quality and reliability ratings as long as the payment for each supplier reflects those values. It is also likely that interdependencies between the attributes exist. For example, the wtp for longer warranty can depend on the product quality. These kind of interactions at the multiattribute level are not ruled out by the model.  $B$  can also express preferences over the number of suppliers and the distribution of quantities using the quantity factor. The restrictions of Theorem 1 begin to arise if  $B$  has particular preferences on the mix of qualities. For example,  $B$  might wish to have some percentage of higher-end monitors (higher resolution, # colors, etc...) and some of lower-end, cheaper ones, for different types of employees. Though not directly supported,  $B$  could specify higher-end suppliers and lower-end suppliers and consider only a partial domain for each. In addition,  $B$  might prefer that all suppliers have the same value for wall-mounting compatibility, so  $B$  will not need to purchase different kinds of such hardware. This cannot be accommodated by mechanisms designed according to Theorem 1, however it can still be expressed by our preference representation scheme as detailed in the next subsection.

### Restricting configuration aggregation

Configuration aggregation, required by Theorem 1, prevents traders from expressing any constraints or preferences on the mix of configurations in their allocation. More specifically, if a trader is aggregating she must be willing to trade any combination of configurations, at prices defined by her wtp per configuration and trading partner. This is quite restrictive, since traders may wish to limit the variance of specific attributes (for example, “items from all suppliers must be delivered at the same time”, or the compatibility issue in the monitors example). Bichler and Kalagnanam (2005) intro-

<sup>2</sup>All proofs are available online.

duced such preferences as *homogeneity constraints*, treating them as hard constraints on acceptable allocations.

Another, more expressive way to describe these types of preferences is to say that configurations have marginal utilities which vary based on the allocation in which they are traded. This can be expressed in the GAI framework by expressing a penalty term which includes the relevant attributes. For example, a trader might have a preference against the aggregation of units of different colors. Assume  $c$  represents the attribute color (with values  $c^1, c^2$ ), and  $a_1, \dots, a_{m-1}$  represent the rest of the attributes. For simplicity we assume there are only two potential sellers,  $s_1, s_2$ . The preference for homogeneity indicates the following:

$$\forall q'_1, q'_2, u([c_1^1, c_2^2, q'_1, q'_2]) < u([c_1^1, q'_1]) + u([c_2^2, q'_2])$$

That means that the utility for an allocation with two configurations of different colors is less than the combined utility of two allocations, each includes only one of the colors. To extend the GAI representation of an aggregating trader to support such expressiveness, we thus have to add another GAI element,  $\hat{c} = \{c_1, c_2, q_1, q_2\}$ . Since the two utility terms on the RHS above are defined over  $\hat{c} \cap t_1$  and  $\hat{c} \cap t_2$  respectively, the explicit GAI decomposition from (3) lends itself to expression of the penalty:

$$u(T_b) = u([t_1]) + u([t_2]) + u([c_1^1, c_2^2, q'_1, q'_2]) - u([c_1^1, q'_1]) + u([c_2^2, q'_2])$$

In this model, the utility of allocation is an additive sum of the utility of trades, as long as the colors are the same.

The key lesson from this discussion is that our preference handling approach—using GAI decomposition over the allocation matrix—flexibly adds complexity only as required by traders' preferences. This framework avoids specifying hard constraints on allocations, and allows for expression of specific non-regularities by adding GAI elements to the model. For the specific conditions of Theorem 1, this framework leads to useful computational and economic results as we demonstrate in the next section.

## MUMA GAI auctions

With the decomposition guaranteed by Theorem 1, it seems natural to design a two-phased mechanism for the multi-unit multiattribute domain: a single-unit multiattribute mechanism to solve *MMP* for all pairs of traders, followed by a multi-unit price-only auction to determine the quantity distribution. In this section we propose such a mechanism for the (one-sided) procurement problem.

### Single-unit GAI auctions

The standard single-unit multiattribute problem poses preference handling challenges of its own due to the exponential nature of the domain. The approach most often found in multiattribute auction literature is to impose the highly restrictive form of fully additive preferences (Bichler 2001). Moreover, to achieve efficiency results most procurement mechanisms in literature require the buyer to provide the sellers with her full utility function prior to bidding. An exception in that sense is the work of Parkes and Kalnaganam

(2005) (PK), which suggests auction types in which bidding is over multiattribute price systems. Their approach is favorable to previous ones in terms of information revelation, but it remains within the limitation of either assuming a fully additive model or exploiting no structural assumptions.

Acknowledging that preferences often exhibit dependencies among attributes, yet may possess much useful structure, we previously proposed (Engel & Wellman 2007) a procurement mechanism that is inspired by PK, but generalizes the underlying preference scheme. The auction employs a price system structured according to a GAI model, and update rules designed to guide traders to their efficient allocations. Structured as it is according to the decomposition implied by Theorem 1 (with its first phase corresponding to *MMP*), its second phase can be amended to treat multi-unit allocation. Following an exposition of the single-unit GAI auction from previous work, we describe in the next section our new extension to the multi-unit case.

Since this is a single-unit auction, traders' valuations are captured by single-unit utility. We make another simplifying assumption throughout this section, that each trader's single-unit preferences are indifferent to the trading partner, meaning we can define  $\forall j, \theta, \mu_b(\theta) \equiv \mu_b(\theta, s_j)$ . There are practical cases in which a buyer wishes to discriminate between suppliers, based on some notion of supplier quality or to compensate for switching costs. Such preferences can often be modeled through attributes in  $A$ .

Preferences of all traders are reflected in a GAI structure  $I_1, \dots, I_g$ , such that  $\bigcup_{r=1, \dots, g} I_r = A$ , resulting in a GAI decomposition of  $\mu(\cdot)$  according to Definition 2.<sup>3</sup> We also require that the GAI structure corresponds to a *GAI-tree* (Gonzales & Perny 2004) or GAI forest, to ensure validity of the compact pricing system. This could potentially increase dimensionality, but does not cause loss of generality.

Our iterative multiattribute auction maintains a GAI pricing structure, through a price  $p^\tau(\cdot)$  corresponding to each sub-configuration of each GAI-tree element. The price of a configuration  $\theta$  at time  $\tau$  is defined as

$$p^\tau(\theta) = \sum_{r=1}^g p^\tau(\theta_r).$$

We use  $\theta_r$  to denote the sub-configuration formed by projecting configuration  $\theta$  to element  $I_r$ .

**Definition 6.** Let  $\alpha$  be an assignment to  $I_r$  and  $\alpha'$  an assignment to  $I_{r'}$ . The sub-configurations  $\alpha$  and  $\alpha'$  are *consistent* if for any attribute  $a_i \in I_r \cap I_{r'}$ ,  $\alpha$  and  $\alpha'$  agree on the value of  $a_i$ . A collection  $\nu$  of sub-configurations is consistent if all pairs  $\alpha, \alpha' \in \nu$  are consistent. The collection is called a *cover* if it contains exactly one sub-configuration  $\alpha_r$  corresponding to each element  $I_r, r \in \{1, \dots, g\}$ .

Note that a consistent cover  $\{\alpha_1, \dots, \alpha_g\}$  represents a full configuration, which we denote by  $(\alpha_1, \dots, \alpha_g)$ .

In GAI auctions, bidders submit *sub-bids* on sub-configurations. The set of *full bids* of a seller contains all consistent covers that can be generated from that seller's

<sup>3</sup>It is possible to adjust the mechanism for the case in which only the buyer's GAI structure is available.

current set of sub-bids. The existence of a full bid over a configuration  $\theta$  represents the seller's willingness to accept the current price  $p^\tau(\theta)$  for supplying  $\theta$ .

To start the auction, the buyer reports her complete valuation (in GAI form) to the mechanism. The initial prices of sub-configurations are set at some level above the buyer's valuations. The auction proceeds in two phases. In the first phase (A), at each round  $\tau$ , bids are collected for current prices, and the auction computes the buyer's preferred set of sub-configurations  $\mathcal{M}^\tau$ . Then, prices are adjusted for the next round, reducing the price of every sub-configuration that has received a sub-bid but is not in  $\mathcal{M}^\tau$ . The price decrement per sub-configuration is by an amount of  $\frac{\epsilon}{g}$ , where  $\epsilon$  can be chosen as desired. The auction switches to phase B when for each active seller  $s_j$  there is a configuration  $\eta_j$  such that all of its sub-configurations are both in the buyer's preferred set and have a sub-bid by  $s_j$ . We show that given truthful, myopic bidding, such a configuration must be approximately efficient. In Phase B  $s_j$  is bidding only on  $\eta_j$ , and it is therefore a single-dimensional auction for a discount term. The seller  $s_j$  that is willing to sell  $\eta_j$  for the highest discount from Phase A prices wins the auction. We showed that this mechanism achieves desirable economic results, as outlined next.

The buyer profit from a configuration  $\theta$  is defined, for a given price system  $p^\tau(\cdot)$ , as

$$\pi_b^\tau(\theta) = \mu_b(\theta) - p^\tau(\theta)$$

and similarly  $\pi_j^\tau(\theta) = p^\tau(\theta) - \mu_j(\theta)$  is the profit of  $s_j$ .

We first formalize the notion of myopic truthful bidding in this context.

**Definition 7.** A seller  $s_j$  is called a *Straightforward Bidder* (SB) if at each round  $\tau$  she bids on  $\mathcal{B}_j^\tau$  as follows: if  $\max_{\theta \in \Theta} \pi_j^\tau(\theta) < 0$ , then  $\mathcal{B}_j^\tau = \emptyset$ . Otherwise let

$$\begin{aligned} \Omega_j^\tau &\subseteq \arg \max_{\theta \in \Theta} \pi_j^\tau(\theta) \\ \mathcal{B}_j^\tau &= \{\theta_r \mid \theta \in \Omega_j^\tau, r \in \{1, \dots, g\}\}. \end{aligned}$$

Crucially, we showed that for an SB seller  $s_j$  whose preferences are structured according to the GAI-tree, any consistent cover generated from her set of sub-bids corresponds to an optimal configuration. Therefore the full bids to which  $s_j$  is responsible are all for configurations which are optimal with respect to current prices. The auction chooses the set of sub-configurations  $\mathcal{M}^\tau$  that corresponds to a set of configurations that are approximately optimal for the buyer at the prices in round  $\tau$ , and it is shown that all of its consistent covers are approximately optimal as well. Therefore, all  $\eta_j$  achieve approximately the maximal buyer profit. The results of Phase A are then shown to be approximately efficient. The quality of the approximation depends on the price decrement  $\epsilon$  and on  $e$ , which denotes the number of edges in the largest connected component of the GAI forest.

We first bound the efficiency result of Phase A for SB sellers. All results assume a truthful buyer.

**Lemma 2.** For SB seller  $s_j$ ,  $\eta_j$  is  $(e + 1)\epsilon$ -efficient.

We then derive overall efficiency for the auction:

**Theorem 3.** Given SB sellers, the auction is  $(e + 2)\epsilon$ -efficient: the surplus of the final allocation is within  $(e + 2)\epsilon$  of the maximal surplus.

Finally, we show epsilon-Nash results for seller strategies.

**Theorem 4.** SB is an  $(3e + 5)\epsilon$  ex-post Nash Equilibrium for sellers in GAI auction. That is, sellers cannot gain more than  $(3e + 5)\epsilon$  by deviating.

On the computational side, we provided an upper bound on the number of rounds needed for Phase A to converge, and asserted that the computational complexity of the tasks associated with Phase A depends only on the number of sub-configurations.

## Multi-unit GAI auctions

Assume traders have MUMA preferences that satisfy the conditions of Theorem 1. In addition, we retain the GAI structure  $I_1, \dots, I_g \subseteq A$  of the single-unit function  $\mu(\cdot)$  as above, and it is independent of the preference structure of the allocation matrix imposed by Definition 3.

We define MUMA GAI auctions as follows. The buyer initially reports her valuation under the structure of Theorem 1, which can be represented using a single-unit valuation function  $\mu(\cdot)$  (decomposed by its GAI structure), trade quantity bounds  $\underline{q}_b$  and  $\bar{q}_b$ , and a quantity factor  $f_b^q(\bar{q})$ . We then execute Phase A, identical to the one defined for the single-unit case. Since  $\mu(\cdot)$  is a multiattribute wtp function, MMP is identical to the notion of seller efficiency, and therefore from Lemma 2 we get that the configurations  $\eta_j$  are approximate MMP solutions for each seller  $s_j$ .

In the single-unit version, Phase B is a price-only ascending clock auction (the price is as the discount term). Similarly, for Phase B in the multi-unit case we use a multi-unit (price-only) auction. For the purpose of the analysis below, we assume a direct revelation VCG mechanism, in which each seller  $s_j$  reveals  $\mu_j(\eta_j)$  and  $\bar{q}_j$  (recall that a seller's quantity factor must be zero when there is only one buyer, and that we require  $\underline{q}_j = 0$ ).

**Theorem 5.** Under the preference restriction of Theorem 1, and assuming SB sellers and a truthful buyer, a MUMA GAI auction with an efficient (price-only) mechanism as Phase B, achieves total surplus which is within  $q\epsilon(e + 1)$  of the efficient allocation, where  $q$  is the number of units allocated by the efficient solution.

If the mechanism used in Phase B guarantees only an approximate solution, within some  $\delta$  of optimum, the proof can be easily adapted to show that the auction would achieve surplus within  $\delta + q\epsilon(e + 1)$  of optimum. For example, the mechanism we apply in the single-unit case is  $\epsilon$ -efficient, resulting in the overall approximate efficiency of Theorem 3.

Finally, we establish incentive properties of the auction based on those of Phase B.

**Theorem 6.** Given a truthful buyer, if the mechanism used in Phase B yields sell-side VCG prices (with respect to its input), then SB is an ex-post  $2q\epsilon(e + 1)$ -Nash Equilibrium for sellers in the MUMA GAI auction. That is, sellers cannot gain more than  $2q\epsilon(e + 1)$  by unilaterally deviating from SB.

As with the previous result, this theorem can be adapted to approximate VCG mechanisms, adjusting the potential gain accordingly. Since the result above assumes nothing on the Phase B mechanism other than efficiency and VCG prices, it applies to any iterative mechanism implementing the VCG outcome. A direct generalization of the single-unit mechanism would be an iterative forward auction over a discount from Phase A prices, (e.g., Ausubel (2004)). However, the problem of non-convex utility on the bid taker (buyer in our case) side (reflected in  $f_b^q(\cdot)$ ), is rarely addressed in the context of iterative mechanisms. We therefore might have to resort to a direct revelation VCG mechanism when a non-trivial buyer quantity factor  $f_b^q(\cdot)$  exists. Indeed, the computational problem associated with the WDP of the VCG mechanism can be intractable, but several useful cases of non-convex utility can be handled using generalized network flow algorithms. These algorithms are shown to perform reasonably well in clearing a double-sided market with up to hundreds of traders (Engel, Wellman, & Lochner 2006).

There are other, well-known drawbacks to a direct revelation approach, for example in the communication burden of full evaluation. However, the advantage we achieve with our mechanism is the reduction of the multiattribute problem to a price-only problem, by the means of using Phase A and Theorem 1. In practice, procurement problems rarely involve more than a few potential suppliers, and the quantity can be discretized to keep the total number of units low. In such case the WDP of Phase B is tractable, but the attributes domain can be extremely large.<sup>4</sup> Therefore, the combination of compact representation for the single-unit preferences and the decoupling of the multiattribute domain from the multi-unit problem are essential in rendering the problem solvable.

### Example

For this example, our buying agent  $B$  procures a large quantity of hard drives, again for the company's internal use.  $B$  is willing to accumulate hard drives of various performance, as long as the payment varies accordingly. Specifically, there are three attributes, with two possible values each:

**Capacity (c)**  $c^1 = 60GB, c^2 = 120GB$ .

**RPM (r)**  $r^1 = 3600rpm, r^2 = 5400rpm$ .

**Warranty (w)**  $w^1 = 3mths, w^2 = 6mths$ .

The GAI structure is  $I_1 = \{r, c\}, I_2 = \{c, w\}$  (note that  $e = 1$ ). The resulting sub-configurations are therefore  $r^1c^1, r^2c^1, \dots$ , and  $c^1w^1, c^2w^1, \dots$

The procurement department considers two potential suppliers,  $s_1$  and  $s_2$ . Since the company wishes, for future reasons, to maintain relationship with both, it has some preference towards dividing the quantity between the two suppliers. For simplicity we divide the large quantity to two equal units, meaning  $\bar{q}^b = 2$ . All traders' utilities exhibit configuration aggregation and linear configuration utility. The

<sup>4</sup>As an example, KLM performed a (single) sourcing event to select a supplier of buttercups for its in-flight service. There were 77 attributes, and only seven potential suppliers, two of which were pre-eliminated (www.negometrix.com).

sellers' lower quantity bound is zero, and their capacity is unlimited in the relevant range (meaning  $\bar{q}^1, \bar{q}^2 \geq 2$ ).

The single-unit utility functions of  $B$ ,  $s_1$ , and  $s_2$  ( $\mu_b, \mu_1, \mu_2$  respectively) are given in Table 1. The buyer's preference towards distribution of supply is expressed through the quantity factor  $f_b^q$  of her utility as follows:

$$f_b^q(q_1, q_2) = \begin{cases} 30 & q_1 = q_2 = 1 \\ 0 & q_i = 2, q_j = 0 \ (i \neq j \in \{1, 2\}) \end{cases}$$

We first analyze Phase A, whose progress is given in Table 2. Let  $\epsilon = 8$ , so the price decrement per sub-configuration is 4. The efficient allocation is  $(s_2, r^2c^1w^2)$ , yielding a surplus of  $145 - 75 = 70$ .  $s_1$ 's efficient configuration is  $r^2c^2w^1$  with surplus of 50. We set high initial prices, 100 for all sub-configurations in  $I_1$  and 75 for  $I_2$ .

Initially, prices of all configurations are equal, leading the sellers to their cheapest configuration ( $r^1c^1w^1$ ) and the buyer's preferred configuration is his most valuable one ( $r^2c^2w^2$ ).  $s_1$  changes his bid to  $r^1c^2w^1$  in round 3, since its current profit is  $175 - 85 = 90$ , compared to  $159 - 75 = 84$  of  $r^1c^1w^1$ . This configuration remains optimal until rounds 8 (though the profit is reduced to  $151 - 85 = 66$ ), in which  $r^2c^2w^1$  becomes optimal too with the same profit.  $s_2$  remains on  $r^1c^1w^1$  until round 4, and alternates between this and  $r^2c^1w^2$  at each round. In round 7,  $r^2c^1w^2$  finally becomes buyer preferred (with buyer profit of  $-14$ ) and in the next round, when  $s_2$  bids  $r^2c^1w^2$  again, both sellers are bidding on configurations that are in the buyer's preferred set. Phase A therefore terminates with  $\eta_1 = r^2c^2w^1$  (which is chosen over  $r^1c^2w^1$  since it yields higher buyer profit) and  $\eta_2 = r^2c^1w^2$ . If using an iterative mechanism, the sellers would compete on a discount from the prices of round 8, and given their valuations  $\pi_1(\eta_1) = 171 - 105 = 66$  and  $\pi_2(\eta_2) = 159 - 75 = 84$ ,  $s_1$  would have dropped out once the discount exceeds 66.

However to take  $f_b^q(\cdot)$  into account, we use a direct VCG mechanism for Phase B. The WDP is as follows:

$$\begin{aligned} GMAP = \max_{q_1, q_2} & (q_1\mu_b(\eta_1) + q_2\mu_b(\eta_2) + f_b^q(q_1, q_2) \\ & - q_1\mu_1(\eta_1) - q_2\mu_2(\eta_2)) = \\ & \max_{q_1, q_2} (f_b^q(q_1, q_2) + 50q_1 + 70q_2) \end{aligned}$$

Since the bonus given by  $f_b^q$  for an even distribution is larger than 20, which is the gain from assigning the second unit to  $s_2$  as well, the solution is  $q_1 = q_2 = 1$ , yielding a surplus of  $30 + 50 + 70 = 150$ . To calculate the price for each seller  $s_j$  we must find the surplus given that  $s_j$  does not participate, meaning that the other trader is allocated both units.

$$GMAP_{-1} = 70q_2 = 140, \quad GMAP_{-2} = 50q_1 = 100$$

This takes into account the loss of the bonus for quantity distribution. The profit of  $s_1$  should be  $150 - 140 = 10$ , and the profit of  $s_2$  should be 50. The final prices are therefore:

$$p(\eta_1) = 105 + 10 = 115, \quad p(\eta_2) = 75 + 50 = 125.$$

### Conclusion

We present a preference handling approach for the domain of multi-unit multiattribute auctions. We show how GAI decomposition can be leveraged to take advantage of structured

	$I_1$				$I_2$			
	$r^1c^1$	$r^2c^1$	$r^1c^2$	$r^2c^2$	$c^1w^1$	$c^2w^1$	$c^1w^2$	$c^2w^2$
$\mu_b$	60	85	70	95	40	60	60	70
$\mu_1$	50	60	55	75	25	30	40	50
$\mu_2$	35	45	65	90	20	40	30	50

Table 1: GAI utility functions for the example domain.  $\mu_b$  represents the  $B$ 's valuation, and  $\mu_1$  and  $\mu_2$  costs of  $s_1$  and  $s_2$ .

	$I_1$				$I_2$			
	$r^1c^1$	$r^2c^1$	$r^1c^2$	$r^2c^2$	$c^1w^1$	$c^2w^1$	$c^1w^2$	$c^2w^2$
1	100	100	100	100	75	75	75	75
	$s_1, s_2$			*	$s_1, s_2$			*
2	96	100	100	100	71	75	75	75
	$s_1, s_2$			*	$s_1, s_2$			*
3	92	100	100	100	67	75	75	75
	$s_2$		$s_1$	*	$s_2$	$s_1$		*
4	88	100	96	100	63	71	75	75
		$s_2$	$s_1$	*		$s_1, *$	$s_2$	*
5	88	96	92	100	63	71	71	75
	$s_2$		$s_1$	*	$s_2$	$s_1, *$		*
6	84	96	88	100	59	71	71	75
		$s_2$	$s_1$	*		$s_1, *$	$s_2$	*
7	84	92	84	100	59	71	67	75
	$s_2$	*	$s_1$	*	$s_2$	$s_1, *$	*	*
8	80	92	80	100	55	71	67	75
		$s_2, *$	$s_1, *$	$s_1, *$		$s_1, *$	$s_2, *$	*

Table 2: Auction progression in phase A. Sell bids are denoted by  $s_1, s_2$  and designation of  $\mathcal{M}^t$  by \*.

MUMA preferences, and its functional form naturally expresses non-regularities in traders' utilities. Specific restrictions allow the decoupling of the multi-unit problem from the multiattribute domain, leading to an auction mechanism that yields approximate efficiency under approximate incentive properties, and takes advantage of GAI structure in both the multiattribute and MUMA domains.

## References

- Ausubel, L. M. 2004. An efficient ascending-bid auction for multiple objects. *American Economic Review* 94(5):1452–1475.
- Bacchus, F., and Grove, A. 1995. Graphical models for preference and utility. In *11th Conference on Uncertainty in Artificial Intelligence*, 3–10.
- Bichler, M., and Kalagnanam, J. 2005. Configurable offers and winner determination in multi-attribute auctions. *European Journal of Operational Research* 160:380–394.
- Bichler, M. 2001. *The Future of e-Markets: Multi-Dimensional Market Mechanisms*. Cambridge Univ. Press.
- Boutilier, C.; Sandholm, T.; and Shields, R. 2004. Eliciting bid taker non-price preferences in (combinatorial) auctions. In *19th Natl. Conf. on Artificial Intelligence*, 204–211.
- Che, Y.-K. 1993. Design competition through multidimensional auctions. *RAND J. of Economics* 24(4):668–680.
- Engel, Y., and Wellman, M. P. 2007. Generalized value decomposition and structured multiattribute auctions. In *8th ACM Conference on Electronic Commerce*, to appear.
- Engel, Y.; Wellman, M. P.; and Lochner, K. M. 2006. Bid expressiveness and clearing algorithms in multiattribute double auctions. In *7th ACM Conference on Electronic Commerce*, 110–119.
- Fishburn, P. C. 1967. Interdependence and additivity in multivariate, unidimensional expected utility theory. *Intl. Economic Review* 8:335–342.
- Gonzales, C., and Perny, P. 2004. GAI networks for utility elicitation. In *9th Intl. Conf. on the Principles of Knowledge Representation and Reasoning*, 224–234.
- Metti, T.; Harlan, R.; Samelson, Q.; et al. 2005. Reinventing the supplier negotiation process at motorola. *Interfaces* 35(1):7–23.
- Parkes, D. C., and Kalagnanam, J. 2005. Models for iterative multiattribute procurement auctions. *Management Science* 51:435–451.
- Sandholm, T., and Suri, S. 2001. Side constraints and non-price attributes in markets. In *IJCAI-01 Workshop on Distributed Constraint Reasoning*.
- Sandholm, T.; Levine, D.; Concordia, M.; et al. 2006. Changing the game in strategic sourcing at Procter & Gamble: Expressive competition enabled by optimization. *Interfaces* 36(1):55–68.