

Enabling assisted strategic negotiations in actual-world procurement scenarios

J. A. Rodríguez-Aguilar, A. Reyes-Moro, M. López-Sánchez, J. Cerquides

ISOCOLab. iSOCO

Intelligent Software Components SA

Edificio Testa c/ Alcalde Barnils, 64/68 A

08190 Sant Cugat del Vallès, Barcelona, Spain

jar@isoco.com, toni@isoco.com, maite@isoco.com, cerquide@isoco.com

Abstract

In the everyday business world, the sourcing process of multiple goods and services usually involves complex negotiations (via telephone, fax, etc) that include discussion of product and service features. Nowadays, this is a high-cost process due to the scarce use of tools that streamline this negotiation process and assist purchasing managers' decision-making. With the advent of Internet-based technologies, it has become feasible the idea of tools enabling low-cost, assisted, fluid, on-line dialogs between buyer enterprises and their providers wherever they are located. Consequently, several commercial systems to support on-line negotiations (e-sourcing tools) have been released. And yet, it is our view that there is still a need for these systems to incorporate effective decision support. This article presents the foundations of *Quotes*, a sourcing commercial application developed by iSOCO that, in addition to cover the whole sequence of sourcing tasks, incorporates decision support facilities based on Artificial Intelligence (AI) techniques that successfully address previous limitations within a single and coherent framework. The paper focuses on the computational realisation of sourcing tasks along with the decision support facilities they require. While supported negotiation events are Request for Quotations/Proposals (RFQs/RFPs) and reverse auctions, decision support facilities include offer generation, offer comparison, and optimal bid set computation (winner determination) in combinatorial negotiations. Additionally, the paper presents a compound of experiences and lessons learned when using *Quotes* for real sourcing processes.

Keywords: negotiation tools, e-procurement, sourcing, combinatorial auctions, multi-attribute auctions, artificial intelligence.

1 Introduction

The sourcing process of multiple goods or services usually involves complex negotiations that include discussion of products' features as well as quality, service and availability issues. Consequently, several commercial systems to support on-line negotiations (e-sourcing tools) have been released. In fact, e-sourcing is becoming an established part of the business landscape (Stephens Inc., 2001).

By industry, the highest penetration (Stephens Inc., 2001) is in chemicals and pharmaceuticals (24 percent) and automobile manufacturing (20 percent), whereas retail, wholesale and distribution, metals and metal products, finance, banking, and accounting are at 4 percent each. Across all industries, e-sourcing adopters are using the technology primarily for MRO goods (71 percent), followed by standard parts (67 percent) and raw materials (57 percent).

However, there is still an enormous challenge confronting users who want to get the maximum value out of e-sourcing. Many companies have realized this value on a transactional level but failed to see it through to the bottom line. The early value of e-sourcing tools has been tactical rather than strategic (Aberdeen Group, 2002), as most adopters got into e-sourcing primarily to negotiate price reductions. Therefore there is still a need for these systems to incorporate effective decision support to enact the focus on strategy.

Traditionally, the core of the sourcing process comprises the following tasks:

- Request for Quotation/Proposal (RFQ/RFP) elaboration;
- Provider selection for RFQ/RFP delivery;
- Offer generation;
- Negotiation through offer-counteroffer interaction or reverse auction; and
- Selection of best offers.

This paper aims at describing how the above-mentioned sourcing tasks are performed by Quotes.

Although several commercial systems to support on-line negotiations have been released, to the best of our knowledge, not a single system can claim to address the full complexity of on-line negotiations. The first generation of sourcing tools merely incorporate single-item, price-quantity reverse auctions mechanisms (see (Ariba 2002) and (Commerce One 2002)). Others only offer basic negotiation capabilities that are usually reduced to a demand-offer matching tool (for example, on-line vertical marketplaces such as MetalSite (MetalSite, 2003)). In general terms, there is a lack of decision support functionalities (decision making in sourcing can involve a few hundred offerings each of which is described by several dozen attributes). Finally, there is a lack of technology support for computationally complex negotiation paradigms, which inhibit the application of interesting models such as combinatorial reverse auctions (Kalagnanam 2003, Sandholm 2002, De Vries et al. 2001) and multi-stage negotiations.

This article presents Quotes (Reyes-Moro et al. 2002) as a solution for strategic sourcing that we believe satisfactorily addresses previous limitations within a single and coherent framework. *Quotes* is a commercial tool that supports all sourcing tasks enumerated above. These tasks are solved by means of different Artificial Intelligence (AI) techniques that might be not theoretically new but that together form an innovative tool that pose an opportunity to proof them useful in real sourcing events.

From the point of view of decision support, we have identified three processes where to apply AI techniques that help buyers and providers in their decision making processes. These three processes have been studied and implemented in *Quotes* with satisfactory results. Although they are thoroughly described along different sections in this paper, next we briefly summarize them:

1. *Automated offer generation.* Providers can translate their business knowledge into *bidding* rules that allow instantaneous and automatic construction of recommended offers. A random neighborhood search algorithm controlled by a rule based system

reasons with these rules in order to construct an initial offer that maximizes both buyer and provider preferences, thus rapidly conducting negotiations to *win-win* situations.

2. *Multi-attribute, multi-item scoring algorithm.* Based on the importance assigned by the buyer to each item attribute in an RFQ (price, quality, delivery time, etc.) and his flexibility to accept offers beyond his preferences, a *fuzzy* offer-matching algorithm scores each offer and ranks it accordingly. In this manner, the buyer can easily discriminate competitive from non-competitive offers. Analogously, providers can benefit from the very same algorithm in order to discriminate incoming RFQs.
3. *Computation of the optimal bid set (winner determination) for multi-item (combinatorial) negotiations.* Given a set of offers for either a multi-item RFQ or an auction, a mixed integer programming algorithm obtains the optimal subsets of offers according to various criteria such as minimization of price, start/finish date, maximization of product quality, etc. Multi-criteria optimisation can be supported by using multi-attribute scoring. Although the used algorithm is not new, the way the problem is formulated can be considered as an important contribution to research in combinatorial negotiations.

The rest of the paper is organised in nine different sections that can be grouped into two main blocks. First six sections correspond to the sourcing tasks enumerated above, whereas the last three sections are dedicated to more general aspects.

- The first block starts with section 2 describing RFQs elaboration. Section 3 introduces the supplier side by introducing the definition of production profiles and provider selection for RFQ delivery. Still on the provider side, Section 4 describes manual and automated (via business rules) offer generation. Next, offer-counteroffer interaction and reverse auctions are presented separately, in sections 5 and 6 respectively. Finally, section 7 details negotiation closing through selection of preferred offers.
- The second block starts with section 8 presenting how different negotiation stages (RFQ/RFP and reverse auctions) can be interleaved in a single sourcing event. Section 9 compiles iSOCO's experiences when organizing real sourcing events with the Quotes tool. Finally, some conclusions and related work are provided in Section 10.

2 Request for quotation (RFQ) elaboration

Quotes supports multi-attribute, multi-item RFQs, enabling the creation of multiple types of RFQs (commodity, catalogue, BOM¹ or group-by). Furthermore it provides the expressiveness needed to cope with multi-criteria negotiation procedures.

Typically a buyer creates an RFQ by sequentially adding items. Each item specifies a product, be it either a good or service. Figure 2-1 shows an *RFQ* composed of several items, a so-called multi-item RFQ. A paradigmatic example of multi-item RFQ occurs at industrial settings. The production plan outlined by some company's ERP (Enterprise Resource Planning) or SCM (Supply Chain Management) application comes in the shape of a list of items to be produced along with the parts required for each product, the so-called bill of material. It is the basis for the buyer to start out multiple sourcing events, each one devoted to the procurement of the parts of each one of the items whose production has been planned out.

The process of including an item is composed of two steps: template selection and RFQ value setting. Next we explain both subprocesses in detail.

¹ BOM: Bill Of Materials

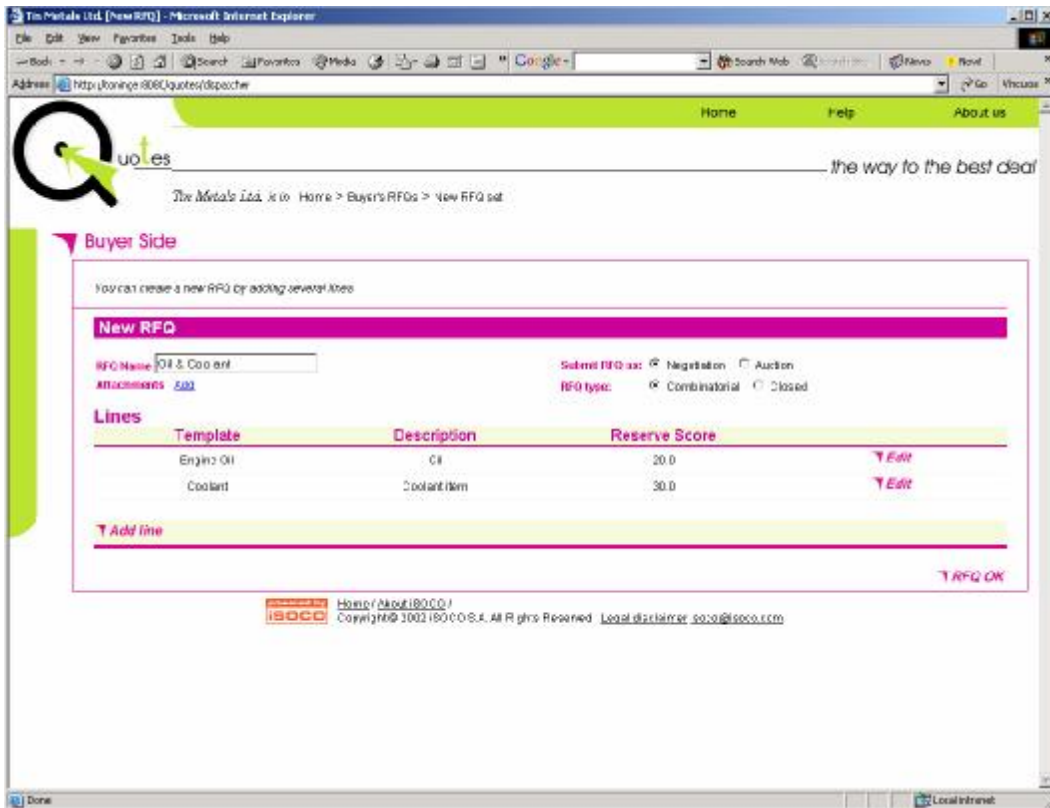


Figure 2-1: Example of a multi-item RFQ

2.1 Item Template selection

The process starts when some buyer creates an RFQ item by selecting a template from a list of product templates. Product templates are previously created by the system administrator and consist of a list of attributes describing the product. Each attribute makes reference to a physical characteristic or negotiable condition or term. Figure 2-2 shows a typical product template for Engine Oil (whose attributes are named Price, Quantity, Viscosity grade, etc.). Each of these attributes is defined in terms of a name and its domain. The name is used as an identifier and provides semantics to the template if a meaningful string is associated. The domain specifies the universe of possible values that an attribute can take on. There are five different types of attributes, namely:

- *Any Number.* For attributes that can take any numerical value. This type is recommended for attributes whose values cannot be constrained (for instance, 'Delivery time' as shown in Figure 2-2).
- *Range of Numbers.* This type is intended for attributes whose possible values belong to a numerical range and it is always associated to a maximum and a minimum value definition. For example, 'Quantity' can only take values between 100 and 1000.
- *Set of Labels.* Non-numerical values can be of type 'set of labels' when there is an associated list of predefined textual values and there is no order among them. For instance, a list of labels ('Conventional', 'Semi-synthetic', and 'Full-synthetic') defines the domain for the (engine oil)'Type' attribute.
- *Ordered Set of Labels.* This type of attribute can be used for quantitative measures. In this case there needs to be an order for associated values. As an example we can think of 'Temperature resistance' degrees such as 'Low', 'Medium', and 'High'.
- *Free Text.* Finally, a free text attribute is added so that both the buyer and provider have the opportunity to add information that is not included in the template (see the 'Additional features' attribute in the example).

Domain specification also includes the measurement units that are associated to numerical attribute values. Hence Quotes enhances each attribute specification with a variety of associated fields that provide both functional flexibility, semantic enrichment, and high expressiveness.

The screenshot shows the 'Administrator Tool' interface for defining an RFQ Template. The template name is 'Engine Oil' and the type is 'Item Template'. The attributes table is as follows:

Name	Description	Type	Values	Units
Price		Range of numbers	0.0 50.0	eur
Quantity		Range of numbers	100.0 1000.0	L
Viscosity grade	N units: winter grade measured at zero F	Ordered set of labels	1, 5 2, 10 3, 15 4, 20	NW
Temperature resistance		Ordered set of labels	1, Low 2, Medium 3, High	
Type	In terms of synthetic composition	Set of labels	Conventional Semi-synthetic Full synthetic	
Delivery time		Any Number		days
Additional features	To be added by buyers	Free text		

Figure 2-2: A product template definition

2.2 RFQ value setting

Once a product template is selected, the buyer must specify desired values for each attribute. First of all the buyer must select a value type per attribute (which in turn depends on its attribute type). Next, the buyer sets the values to be requested per attribute along with its relative importance with respect to the rest of attributes. Figure 2-3 shows the interface available for buyers at this aim.

In this manner, for an *Any Number* attribute a buyer can specify its desired value as either a single number or as a range of numbers satisfying his needs. In the last case, three values are required: a minimum, a maximum, and a preference slope indicating which values are most preferred within the range. The slope can take on three different values:

- *Flat*. All values within the range are equally preferred.
- *More Is Better (MIB)* indicates that higher values are more preferred than lower ones.
- *Less Is Better (LIB)*. On the contrary to MIB, the buyer chooses an LIB slope when smaller values are preferred.

The same kind of single / range value specification can be used for the *Range of Numbers* and *Ordered Set of Labels* types. The only difference is that attribute values must belong to the domain defined by the corresponding attribute template. As an example, consider specifying

that we are prepared to accept any value for attribute 'Viscosity grade' that is at least '15W' but taking into account that the more viscosity the better.

As to the *Set of Labels* type, the buyer is enabled to choose one value (label) or a subset of values. Notice though that since no order is associated, there is no need for defining any range or slope.

Finally, a buyer assigns some value to a *Free Text* type by filling out a text field. In the example, it specifies '50L Barrels' as an additional preference.

Name	Value Type	Value	Units	Must Have	Importance	Reveal
Price	Interval	0.0 10.0 40.0 ... 300 Slope: Less is better	EUR	<input type="checkbox"/>	(5 circles)	<input checked="" type="checkbox"/>
Quantity	Single	0.00.0 400.0 ... 1000.0		<input type="checkbox"/>	(5 circles)	<input checked="" type="checkbox"/>
Viscosity grade	Interval	1E 20 Slope: More is better	W	<input type="checkbox"/>	(5 circles)	<input type="checkbox"/>
Temperature resistance	Interval	Medium High Slope: Any is good		<input checked="" type="checkbox"/>	(5 circles)	<input checked="" type="checkbox"/>
Type	Single	<input type="radio"/> Conventional <input type="radio"/> Semi-synthetic <input checked="" type="radio"/> Full synthetic		<input type="checkbox"/>	(5 circles)	<input checked="" type="checkbox"/>
Delivery time	Interval	1 7.0 21.0 ... Slope: Any is good	Days	<input checked="" type="checkbox"/>	(5 circles)	<input type="checkbox"/>
Additional features	Free text	50L Barrels		<input type="checkbox"/>	(5 circles)	<input type="checkbox"/>

Figure 2-3 : RFQ value specification.

The buyer has also the capability of defining the importance that each attribute is expected to have during the negotiation process. Thus, the buyer assigns a value ranging from *irrelevant* to *extremely important* (as for the 'Temperature resistance' attribute). Additionally, if it is compulsory that offers submitted by providers satisfy the requested values, the buyer must then tick the *must-have* checkbox (as for the 'Temperature resistance' and 'Delivery time' attributes). The way Quotes exploits preferences and importance values is latter explained.

An interesting issue that is beyond the scope of this paper is the problem of *preference elicitation*, i.e. discover and/or quantify the buyer's preferences over a product configuration (Bichler 2002, Peyman 2001).

Finally, the buyer assigns a reserve score, a threshold value, to the RFQ. Thus offers whose degree of matching with the RFQ fall below the reserve score are filtered out (subsection 7.1 fully explains how offers are scored in relation to RFQs). That is, depending on the matching score value, Quotes automatically rejects offers that unsatisfactorily match buyers' requirements.

3 Automated provider selection

3.1 Providers' preferences and capabilities: production profiles

While buyers need to specify their product requirements in terms of negotiable attributes, providers can analogously do the same regarding their product capabilities and their preferences over incoming RFQs. Thus each provider is capable of declaring the so-called production profiles thoroughly detailed lists of products he is capable of providing. Figure 3-1 shows a production profile corresponding to the product specified in the template shown in Figure 2-2.

The screenshot shows a web browser window titled 'All Oil Ltd's Business Preferences - Microsoft Internet Explorer'. The address bar shows 'http://tonings:8080/quotes/idsupplier#attributes'. The page content is a form titled 'Change your Conditions and preferences for this filter'. It includes a 'Product filter' section with a text box containing 'Oil product profile'. Below this is a table of attributes with columns: Attribute Name, Allowed values, Likes/Bothers, Preferred Type, Preferred Values, and Importance. The attributes listed are Price, Quantity, Viscosity grade, Temperature resistance, Type, and Delivery time. Each attribute has specific input fields and options for configuration.

Attribute Name	Allowed values	Likes/Bothers	Preferred Type	Preferred Values	Importance
Price	Min: 20.0, Max: 40.0	<input checked="" type="checkbox"/>	Interval	10.0, 30.0, 40.0, 50.0 Slope: More is better	## ## ## #
Quantity	Min: 300.0, Max: 1000.0	<input checked="" type="checkbox"/>	Interval	100.0, 500.0, 1000.0, ... 1000.0 Slope: Any is good	## ## ## #
Viscosity grade	Min: 10, Max: 20	<input checked="" type="checkbox"/>	Single	<input type="radio"/> 5 <input checked="" type="radio"/> 10 <input type="radio"/> 15 <input type="radio"/> 20	## ## ## #
Temperature resistance	Min: Medium, Max: High	<input type="checkbox"/>	Single	<input type="radio"/> Low <input checked="" type="radio"/> Medium <input type="radio"/> High	## ## ## #
Type	<input type="checkbox"/> Conventional <input checked="" type="checkbox"/> Semi-synthetic <input checked="" type="checkbox"/> Full synthetic	<input checked="" type="checkbox"/>	Soft or labels	<input type="checkbox"/> Conventional <input checked="" type="checkbox"/> Semi-synthetic <input checked="" type="checkbox"/> Full synthetic	## ## ## #
Delivery time	Min: 7.0, Max: 30.0	<input type="checkbox"/>	Single	30	## ## ## #

Figure 3-1: A sample of production profile

As discussed above, Quotes describes a product in terms of its negotiable attributes. Consequently, production profiles need to be specified for each attribute.

When specifying their production profiles, providers are requested to specify both their production capabilities and their selling preferences:

- Production capabilities (shown as 'Allowed values' in Figure 3-1) describe the products that can be offered. They are specified through the range of attribute values that can be provided. For example, suppose that some provider needs a minimum of seven days to deliver his goods or that he can only provide a minimum 'Viscosity grade'. Since production capabilities determine which product demands the provider can actually accept, Quotes uses this information for RFQ delivery.
- Selling preferences allow a provider to state which requests he may favour. This is a way of reducing the high cost involved in analysing large collections of RFQs. A provider might be interested in quickly identifying requests for large volumes or for a specific product model, while being less interested in requests for discontinued

products. As Figure 3-1 shows, selling preferences are specified by means of a preferred values column and, once again, an associated importance. With this information, Quotes is able to aid the provider when deciding which RFQs to consider first (ordering RFQs is detailed in subsection 7.1.3 and is based on scoring computation).

3.2 Provider selection

Given a brand-new generated RFQ to request for a set of goods, provider selection is the automatic process that guarantees that only those providers supplying some of the required goods do receive the RFQ. In addition, providers will see customized views of these RFQs, so that the *RFQ view* for a given provider exclusively contains those items that he or she is capable of providing.

Provider selection consists of two filtering steps. It has been already described how buyers and providers have interfaces that allow them to describe their necessities and capabilities based on commonly-shared product templates. The first filtering process solely requires identifying those providers offering products specified with the same template than the product the buyer is requesting. For every identified provider, the second filtering process focuses on attribute values. Internally, this commonly-shared language provided by templates allows the matching of each provider's capabilities to the buyer requested values given that they belong to the same domain. As mentioned above, when a buyer specifies his requirements by assigning values to each attribute in a template, these values can be single values, sets of values or intervals. Additionally, a buyer can set to true a *Must Have* flag to indicate that the requested attribute value is mandatory. That is, for example, if the buyer requests engine oil with 'Temperature resistance' attribute values between 'Medium' and 'High' and selects *Must Have* (as shown in Figure 2-3), only those providers that have the capability to offer this item with values within this interval will see the item in their *RFQ view*. On the contrary, if it is not the case that this *Must Have* flag is active, then a provider producing engine oil whose 'Temperature resistance' is 'Low' ought to receive the request, since that's what negotiation is all about. Analogously to the *Must Have* flag, providers specify the attributes of the items they produce with an associated *Will Listen Broader* flag. This flag specifies whether a provider is willing to receive requests of items whose attribute values fall outside his capabilities. In this manner, in the previous example, the provider producing low temperature resistance oil would receive the request only if its *Will Listen Broader* flag was active (see third column in Figure 3-1). This flag allows providers to customize their RFQs filtering.

Flag value combination (*Must Have* on the buyer side and *Will Listen Broader* on the provider side) are taken into account before requests are sent to providers. Each pair of values defines a case for which an interval or set intersection check must be performed over requested and produced attribute values so that *RFQ views* will be defined for products whose attributes accomplish these tests.

Table 3-1 presents this combination of values and checks, being:

- RFQAttribute.MustHave the *Must Have* flag value on the buyer side
- ProviderAttribute.WillListen the *Will Listen Broader* flag value on the provider side
- RFQAttr the set of buyer's attribute value preferences
- ProvAttr the set of provider's attribute value capabilities

	RFQAttribute.MustHave=1	RFQAttribute.MustHave=0
ProviderAttribute.WillListen=1	RFQAttr \subseteq ProvAttr?	No test required
ProviderAttribute.WillListen=0	RFQAttr \subseteq ProvAttr?	RFQAttr \subseteq ProvAttr?

Table 3-1: Flag value combination for capability matching.

4 Offer generation

Once an RFQ has been delivered to selected providers, those providers can then generate their first offers. Quotes allows offers to be elaborated either manually by providers or automatically by the system. Nevertheless, since providers in real sourcing events prefer to have complete control over their offering, the automated offer generation functionality limits to composing indicative offers (as opposed to firm offers).

4.1 Manual offer elaboration

Figure 4-1 shows an RFQ item description as seen by a provider. It contains its attributes, their importance (as horizontal bars beside each attribute name), their preferred values and, when these are intervals, the corresponding preference interval slope (this is shown through bold figures representing most preferred values, such as 10.0 euro in the price interval). There, a provider can specify offered values along the New Offer column. The only restriction being that they must belong to the template specifications (attribute value domains appear as intervals next to each textbox in the New Offer column or as the set of values that can be selected within a combo box). Additionally, if no *Must Have* flag is active, offered attribute values can lay outside preferred intervals.

Notice also that, whilst an RFQ item represents a subspace of preferred attribute values within the space defined by the template, offers are restricted to represent points in the very same space. The reason to constrain providers to single-valued offers is mainly semantic, since assigning more than one value to each attribute could be ambiguous in representing interdependencies and could in fact be interpreted as different alternative offers. As an example, consider offering two values of 'Price' and 'Viscosity grade' attributes, this could mean that one viscosity grade is more expensive than the other, or that there are two prices for each viscosity grade. In this manner, single-valued attribute offers do not require to disambiguate all possible value combinations by increasing the complexity of the offer specification language: providers specify attribute value alternatives (disjunction) as separated offers.

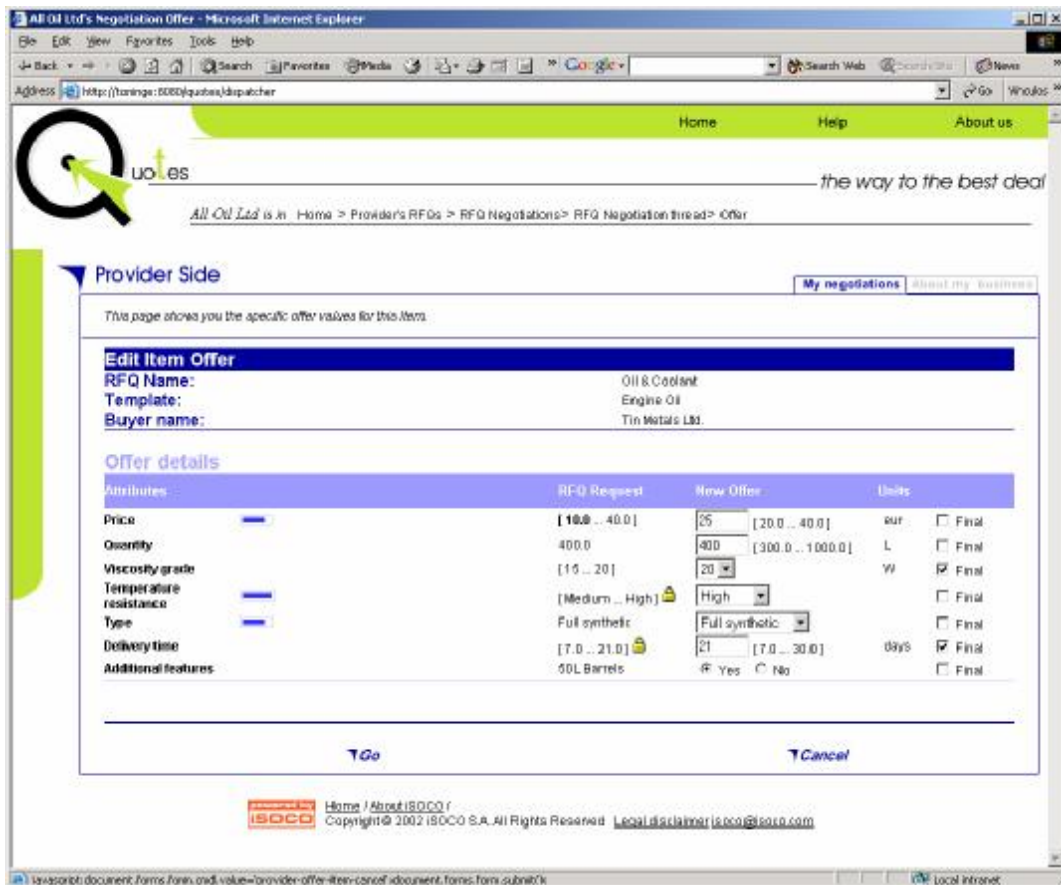


Figure 4-1. Offer edition for an RFQ item

4.2 Automated provider response

The system requires providers to specify their business rules so that they can be subsequently applied to automatically generate offers. Next we detail how these two steps are undertaken.

4.2.1 Business rules as bidding rules

Beyond declaring attribute capabilities and preferences, Quotes allows each provider to declare his own business rules in the form of bidding rules.

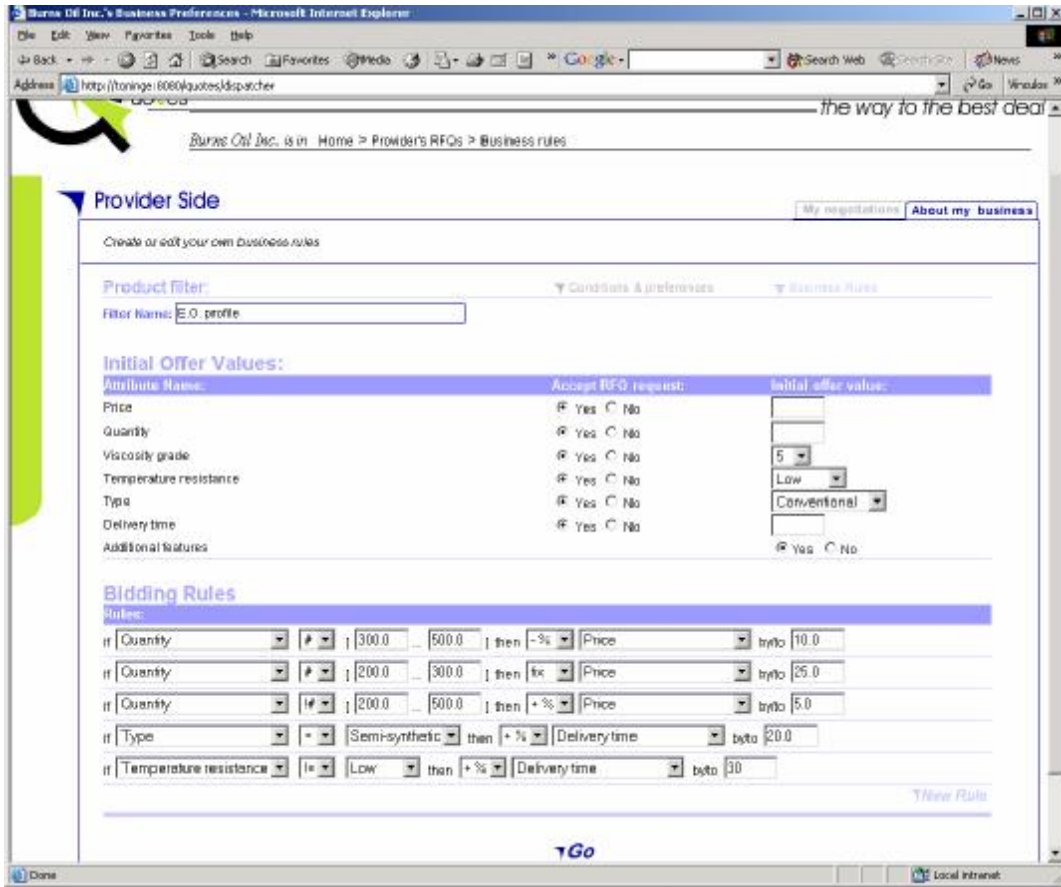


Figure 4-2: Bidding rules sample

A bidding rule is an *if-then* rule that checks and changes the value of one or several attributes. Examples of rules include discount per volume, additional charges for express delivery, no delivery charge when a minimum price is offered, etc. Figure 4-2 shows an example for Engine Oil. The general syntax to define bidding rules is shown in Figure 4-3.

```

RULE: if CONDITION then ACTION;
CONDITION: attribute_name1 [=, !=] value1 |
           attribute_name1 [#,!#] (min_value, max_value) |
           (CONDITION) |
           CONDITION [and, or] CONDITION;
ACTION: attribute_name2 [fix, +%, -%] value2

```

Figure 4-3: Bidding rule syntax

Operations for checks and changes can be selected among a list of available operations that depend on the chosen attribute type:

- *Condition operators* are: =, !=, #, !#. Equality (=) and inequality (!=) operators can be chosen for all types of attributes (both numerical and sets of labels) and require a value to be compared with. On the contrary, membership (#) and non-membership (!#) operators only apply to types that allow to specify intervals (that is, *any number, range of numbers, and ordered sets of labels*).
- *Action operators* are functions to be applied to the attribute that has been chosen to be affected. These are: assignment (fix), which can be applied to all attribute types, and

percentage increment (+%) and decrement (-%), which are only applied to numerical types.

Considering this syntax, the sample rules in Figure 4-2 set the price depending on the quantity (three first rules), increase the delivery time for semi-synthetic oil types (fourth rule) and for medium and high temperature resistant oils (fifth rule).

Notice that the free text type is not supported because it is a natural language value whose automatic treatment goes beyond the scope of this application. However, the manual setting of a boolean value associated to free text types specifies if the provider meets the requested attribute value or not.

4.2.2 Automated offer generation

In this section we describe the module of *Quotes* that is in charge of generating indicative offers on behalf of providers exploiting the above-mentioned bidding rules. Such process takes place immediately after provider selection. Its objective is to build a complete offer (where each attribute is assigned some value). Moreover, the algorithm pursues to build the *best* complete offer in terms of either the buyer's preferences or the provider's preferences or both. For the sake of clarity, this section assumes an offer to be composed of a single item (that is, it is only offering one product), so that 'offer' is used instead of 'offer item'. For the general case of offering more than one product, the process described below will be repeated for each item.

The implemented algorithm is a variation of the well-known hill-climbing random neighbourhood search procedure. We start with an incomplete offer O as the candidate solution and then enter into an iteration phase. At each iteration, we generate a neighbouring offer O' of O to which we apply the provider's business rules (which have been already explained in the pervious subsection). If the new solution O' is better than the candidate offer O we accept O' as the current solution. The process continues until the termination criterion (no improvement for the last k iterations) is reached.

In order to explain completely the optimisation procedures, the following definitions are in place.

Definition 1: Space of solutions.

We define an offer as a tuple o of the form $\langle o_1, o_2, \dots, o_n \rangle$ where $o_i \in (A_i^p \cap A_i^b) \cup \emptyset$, $i=1, \dots, n$, being:

- n the number of negotiable attributes of the offered item;
- A_i^p the set of values produced by provider p for the i -th attribute (a_i);
- A_i^b the set of values ask by buyer b for the i -th attribute (a_i);
- the symbol \emptyset denotes unassigned value.

In other words, *an offer* o sets a value for each requested attribute (i.e., o_i is the value offered for attribute a_i), provided that such value is within both provider's capabilities and buyer's acceptable values.

We say that an offer o is *complete* if $o_i \neq \emptyset \ \forall i=1, \dots, n$ and *partially complete* if $\exists i$ such that $o_i = \emptyset$. Additionally, we say that an offer o is *more partially complete* than o' (denoted as $o \mathbf{f} o'$) if o has less *unassigned* values than o' .

The evaluation function defines the objective of the optimisation process in terms of determining if o' is better than o . In our case the objective is twofold: obtain a complete offer that optimises a target function.

Definition 2: Objective function.

The objective function $c: O \rightarrow \Re$ is defined over the set of offers O as

$$c(o) = w^p \cdot \sqrt{s^p(o)} + w^b \cdot \sqrt{s^b(o)}$$

where

- $s^p: O \rightarrow [0..100]$ is the scoring function based on the preferences of provider p (as detailed in subsection 7.1).

- $s^b: O \rightarrow [0..100]$ is the scoring function based on the preferences of buyer b .
- w^p and w^b are weighting factors so that $w^p + w^b = 1$ and $w^p \geq 0$, $w^b \geq 0$

Such evaluation function tries to favour high scorings while penalizing big differences between the buyer's and the provider's revenues. In other words, it prefers a 50 - 50 rather than an 80 - 20 thus seeking win-win situations (and this is desirable because if a provider generates an offer that best suites his/her preferences without taking into account the buyer's ones, it will be likely rejected by the buyer). Finally, w^p and w^b can be used to tune the objective function (for example, to focus more on the buyer's preferences).

Optimisation algorithm

Informally, we say that an offer o' is *better* than an offer o if o' is more *partially complete* than o and $c(o') > c(o)$. Consequently, we only favour solutions that are closer to be complete and that improve the existing candidate solution. Figure 4-4 shows the pseudo-code for the optimisation algorithm.

```

o = <∅, ∅, ..., ∅>
While termination_criterion not reached
    randomly select an attribute  $a_i$  in the condition of a bidding rule.
    randomly obtain a new value  $o'_i$  for the chosen attribute.
    obtain  $o'$  as the result of executing the set of bidding rules over  $\langle o_1, \dots, o'_i, \dots, o_n \rangle$ 
    if  $o' > o$  then  $o = o'$ 
end_while
If  $O$  is complete return with success
else return with failure

```

Figure 4-4. Optimisation offer algorithm

If the process ends with success, a complete *offer* is obtained and Quotes automatically submits it as an indicative offer. Assuming $w^p = w^b = 1$ the offer is likely to be close to a win-win agreement and thus reducing the number of buyer-provider interactions. However the negotiation process may progress further since the submitted offer is set as indicative.

4.2.3 Example

Consider a buyer request consisting of 10 units (non-mandatory) of memory SIMMS of exactly 128Mb at 133Mhz. Furthermore, the buyer is not willing to pay more than 10000. Additionally, we assume there is a provider of SIMM Modules whose bidding rules are shown in Figure 4-5.

Bidding Rules

Rules:

if	Size	=	64 Mb	then	fix	Price	by/to	5000.0
if	Size	=	128 Mb	then	fix	Price	by/to	8000.0
if	Size	=	256 Mb	then	fix	Price	by/to	15000.0
if	Speed	=	133 Mhz	then	+	%	Price	by/to 30.0
if	Quantity	#	[11.0 ... 100000]	then	-	%	Price	by/to 10.0

New Rule

Go

Figure 4-5: Bidding rules example

Table 4-1 shows the evolution of an offer through the optimisation algorithm execution. Since we assume the provider has not specified any preferences (and thus it is not possible to compute his scoring), we only show the evolution of the buyer's score.

Iteration	Price	Quantity	Size	Speed	Package	$s^b(o)$
0	∅	∅	∅	∅	∅	--
1	∅	∅	∅	133	∅	--
2	∅	∅	∅	133	Branded	--
3	∅	8	∅	133	Branded	--
5	10400	8	128	133	Branded	35%
12	9960	19	128	133	Branded	66%
36	9960	11	128	133	Branded	80%

Table 4-1: Evolution of an automatically generated offer

The algorithm quickly assigns values for speed, package and price. Speed and price only admit one value each, and the value selected for Package has no effect on the buyer's score. At fifth iteration, the algorithm has managed to obtain a complete offer. At 12th iteration, a random move of quantity enables the application of a volume-based discount rule (last bidding rule in Figure 4-5). This improves the buyer's score since it fixes the price below 10000. From this point, the only allowed movement is to decrease the number of units to match as much as possible the quantity required by the buyer. The algorithm successfully terminates with the offer <9960,11,128,133,Branded>, which guarantees the price to be below 10000 offering only an additional unit.

5 Negotiation

So far potential providers have been requested for offers and even some of them have already submitted either manual offers or automatically-generated indicative offers. The process now enters into a negotiation phase. Negotiation is conducted through multiple, simultaneous, structured dialogs. Each dialog is established between the buyer and a single provider and it is ruled by a negotiation protocol. Buyer actions can be: offer acceptance, offer rejection, counter-offer submission, and request for firm (offer). Provider actions are limited to the submission of either firm or indicative offers. Nevertheless, a provider can negotiate for an RFQ both by submitting offers sequentially (following a negotiation dialog) or in parallel (offering different alternatives to the buyer).

The following sequence illustrates a typical negotiation.

1. The buyer submits an RFQ asking for *service1* and *service2*.
2. *Quotes* identifies a potential provider and automatically constructs two indicative offers on his behalf based on his bidding rules: *offer1* for *service1* and *offer2* for *service2*.
3. The buyer evaluates *offer1* and submits a counter-offer asking for lowering the price.
4. The provider responds with an extension of *offer1* so that it also includes an offer for *service2*. In other words, he is accepting a price reduction provided that the buyer acquires both *service1* and *service2*.
5. The buyer rejects *offer2*.
6. The buyer evaluates the modified *offer1'*, agrees with it and requests a firm offer.
7. The provider concedes the firm offer.
8. The buyer accepts the firm offer and therefore the negotiation successfully finalises.

Formally, negotiation can be described by means of a finite state automaton where messages between a buyer and a provider define the transition between negotiation states. Next Figure 5-1 depicts it.

Notice that the buyer holds a one-to-many negotiation with providers. *Quotes* allows that multiple dialogs take place in parallel between the buyer and the providers competing for some RFQ. And at the same time, each provider is allowed to simultaneously hold several dialogs with the buyer whenever each dialog corresponds to a different offer.

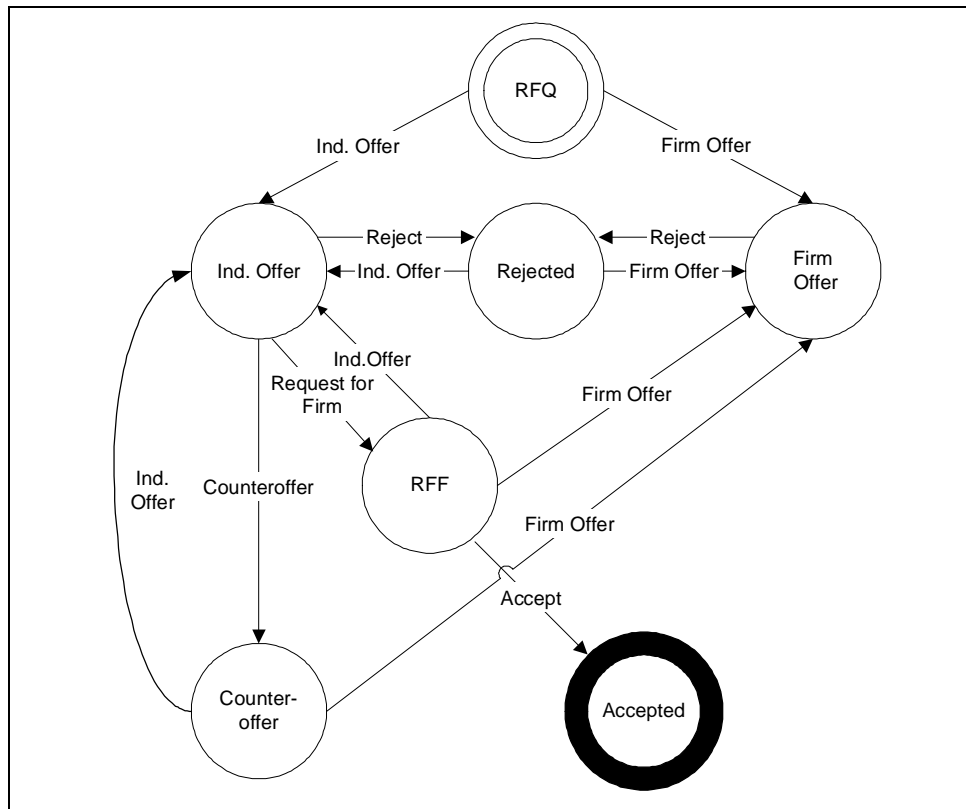


Figure 5-1: Negotiation finite state automaton.

As we have previously stated, negotiation assumes an RFQ has been sent, so negotiation formally starts with a provider sending an offer. This initial state switches either to an *Indicative Offer* state (reached when an “Ind. Offer” message is sent) or to a *Firm Offer* state. Intermediate states are Counter-offer, RFF (Request For Firm), and Rejected.

Negotiation ends when the buyer accepts a firm offer. This corresponds to an *Accepted* state (in black). Another final state is *Failure*, which can be reached from any other state (although it has been omitted for the sake of clarity). In addition to this, neither the buyer nor provider are directly responsible for generating the transition to this state: Quotes causes it as a consequence of an internal action (such as a system exception, a timeout, etc.).

6 Auctions

In order to start an auction, an RFQ is required together with a selection of qualified providers. Only selected providers will be invited to participate in a buyer-customized auction event. Auctions in Quotes include several parameters (Sandholm, 2000) (most of them are shown in Figure 6-1):

- When to clear the auction. Available options are: to be cleared by buyer, when a specific time is reached (thus, an ending date must be provided), or when no bids have been received for a specified time.
- Begin date: when to launch the auction. Prior to the auction start, its terms are accessible to invited providers.
- First bid unconstrained allows a provider to bid for the first time with an offer that does not necessarily overbid the best bid. As a consequence the winner remains the same but auction competitiveness is increased.
- What information is revealed to bidders. Three different levels of information visibility concerning providers' identities can be established: none, nickname or full identity.
- What bidding information is revealed to bidders while the auction stage (highest bid, all bids, or none).
- Ranking revelation. If enabled, each bidder can visualise the relative position of his bids.
- Voyeur providers. Whether providers that have not submitted any bid can see information about the auction evolution or not.
- Maximum number of auction extensions and time per extension.
- Extension detection time. Extensions are launched when bids are submitted within this period of time.

Additionally, bidding constraints can be imposed on providers, lots, items, or attributes:

- Constraints on providers affect bidding capabilities for certain bidders, as for example, different starting values for attributes, fixed bids for specific items, etc.
- Constraints on lots and items fix the maximum/minimum bid increment/decrement.
- Constraints on attributes specify starting and target values as well as increment/decrement constraints.

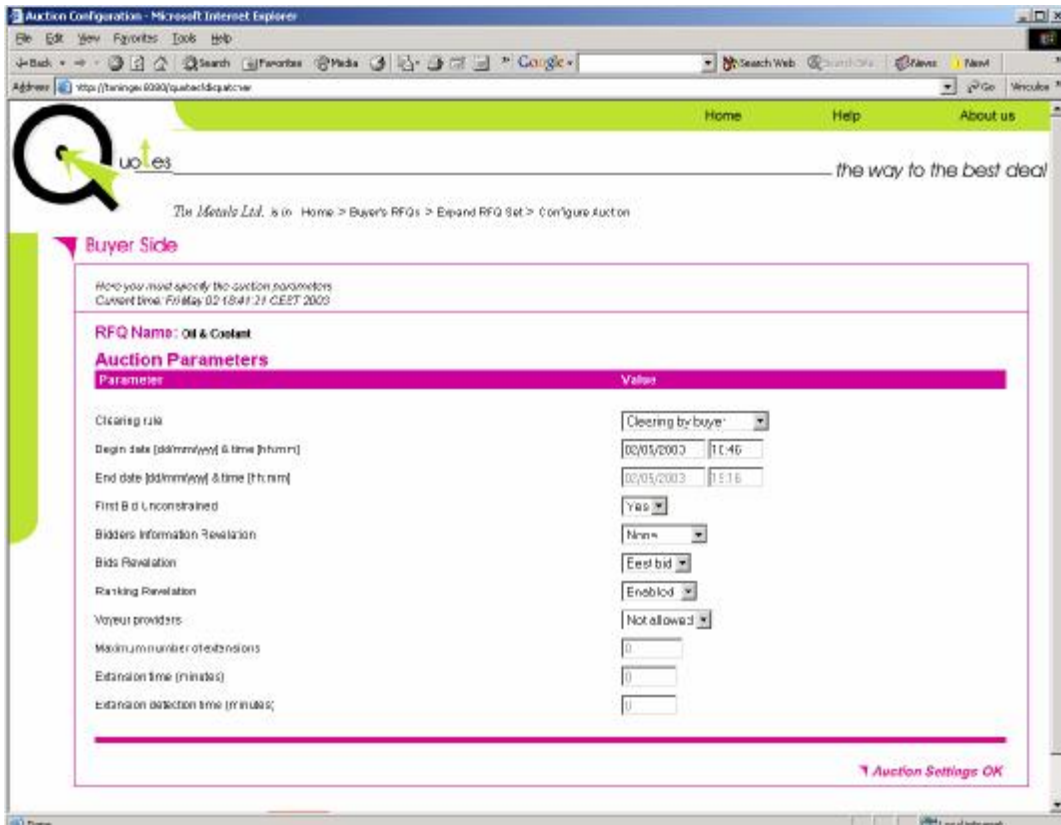


Figure 6-1. Auction parameters' setting

In the most general case, Quotes supports combinatorial multi-attribute reverse auctions (Sandholm et al. 2002). This allows providers to directly bid for bundles of items. They are convenient for providers that have non-additive values for bundles of items (providers' offers for bundles of items may be better than for their separate items). Furthermore, they allow providers to express complementarities (interdependencies) over the requested items to avoid the risk of being awarded incomplete bundles. Notice also that providers are allowed to place multiple bids for bundles of items. Figure 6-2 and Figure 6-3 depict how buyers and providers visualise the very same combinatorial auction.

Tin Metals Ltd. Auction View - Microsoft Internet Explorer

Address: http://tinmetals.com/qauction/auctionview

Tin Metals Ltd. is in Home > Buyer's RFQs > RFQ Auction

Buyer Side

Have you saved the RFQ auction state with each line details
Current time: Fri May 02 16:43:03 CEST 2003

Auction PUBLISHED

Automatic page refresh

RFQ Name: **Oil & Coalbert**

Auction Settings: [View](#)

Invited Providers: [View](#)

Begin date: **May 2, 2003 6:46:00 PM**

Time to begin: **0 d, 00:02:07**

Lines

Line	Template	Line name	Status				
1	045	Oil		Edit			
ID	Provider	Ranking	Bid Type	Score	Reception Time	Bid state	
0247	All Oil Ltd	1	Composed	<div style="width: 100%;"></div>	May 2, 2003 3:38:01 PM		
2	046	Coalbert Item		Edit			
Compare ID	Provider	Ranking	Bid Type	Score	Reception Time	Bid state	
<input type="checkbox"/> 0247	All Oil Ltd	1	Composed	<div style="width: 100%;"></div>	May 2, 2003 6:20:01 PM		
<input type="checkbox"/> 0243	Burns Oil Inc	2	Single	<div style="width: 100%;"></div>	May 2, 2003 6:38:28 PM		
Compare							

Figure 6-2. Buyer façade of an ongoing auction

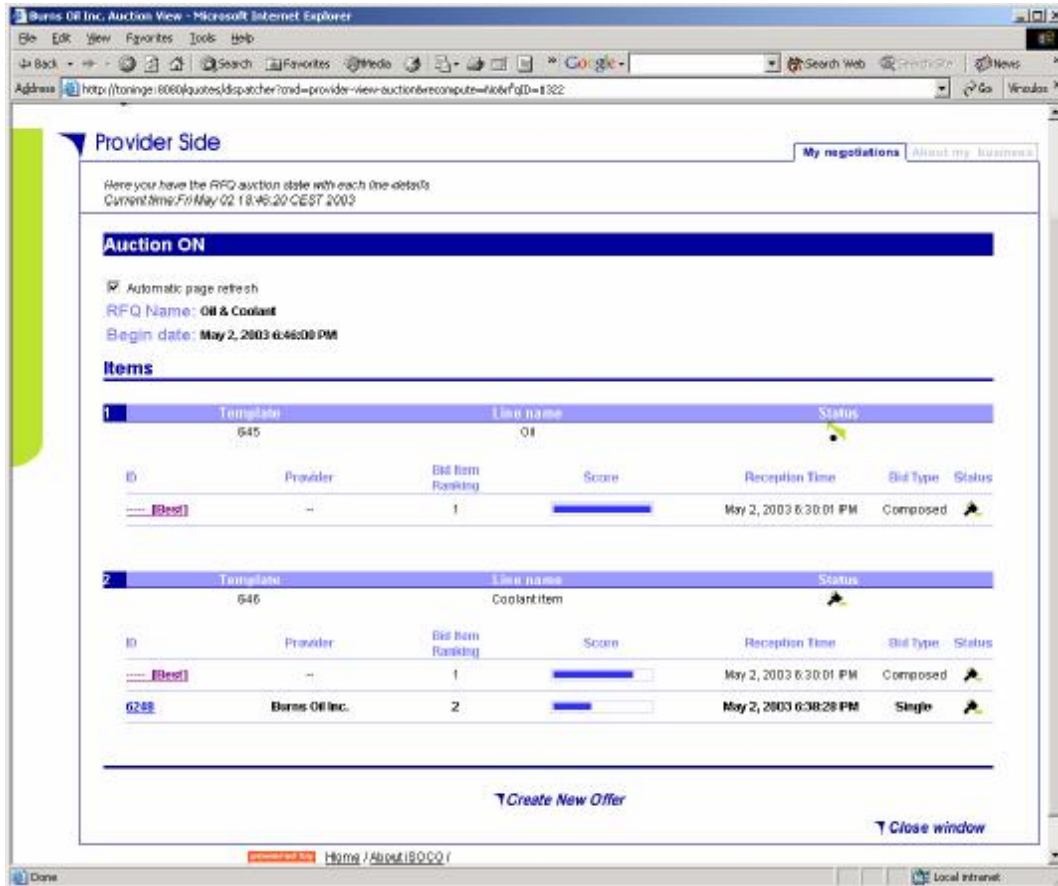


Figure 6-3. Provider façade of an ongoing auction

7 Selection of best offers

Ideally, both negotiations and reverse auctions should finish with the award of items to the best offers. Nevertheless, when dealing with complex (multi-item, multi-attribute) negotiation events, this is not a straight forward task. In order to assist buyers, Quotes provides a decision support module (the so called RFQ-offer matching module) for offer and bid assignment endowed with ranking and comparing facilities. They are based on scoring criteria fully detailed below.

Furthermore, Quotes incorporates an additional decision support module that determines the best offers in combinatorial negotiations whose implementation is presented in subsection 7.2.

7.1 Scoring criteria: Fuzzy matching

Quotes provides both buyers and providers with a fuzzy matching module that allow them to score the negotiation messages (RFQs and offers) they receive based on their own preferences. In this manner, a buyer can order incoming offers from different providers in the same way that a provider can order incoming RFQs from different buyers. This is specially useful when dealing with a large number of negotiation messages because the more interesting a message the earlier it should be identified and answered. And the sense of interest is extracted from buyers' and providers' preferences.

Most commercial offer selection tools are based on simple implementations of Multi attribute utility theory (MAUT, Keeny and Raiffa 1993). We extend these techniques by incorporating fuzzy functions in the RFQ-offer matching module (see Ribeiro 1996). At this aim, we firstly represent both requested and offered attribute values as fuzzy functions. Secondly, this pair of fuzzy functions are combined and defuzzified (by computing the supremum of their intersection) in order to obtain a scoring (a degree of matching) at attribute level. These crisp values are then weighted with the importance of each attribute so that the scoring for an item is

obtained. Finally, the scorings for all items in a message are aggregated to yield a total scoring value.

After the work by Baas and Kwakernaak (1977) fuzzy functions have been mostly applied as triangular fuzzy numbers representing preferences. On the other hand, preferences over continuous attributes can be modelled by linear functions in the $[0, 1]$ interval. We go further in three aspects: we parametrize fuzzy functions' support (positive values), we model interval preferences with trapezoidal fuzzy functions, and we allow values in the central part of the preferred intervals to increase or decrease linearly. The three extensions are determined based on users' preferences.

7.1.1 Fuzzy functions

As we have previously seen, both buyers and providers define their preferences. On the one hand, buyers specify their preferences when assigning values to RFQ item attributes (see subsection 2.2). On the other hand, providers specify their preferences when defining preferred values in their product profiles (see subsection 3.1). Internally, these preferences are represented as fuzzy functions.

A remarkable feature of fuzzy logic (Godo et al., 1993) is its ability to handle the concept of relative truth of one proposition "x is P" through the specification of a membership function that represents predicate P. In our case, we can see a preference as a predicate, and the degree of truth of the proposition as the degree of the preference satisfaction of an offered value x. For example, consider a domain of four quality values ('low', 'medium', 'high', and 'luxurious'), a buyer requesting 'luxurious' quality, and a provider offering 'high' quality. In this case we cannot state that the provider satisfies completely the buyer's preference, but, since 'high' is close to 'luxurious', satisfaction should neither be zero.

In Quotes, this degree of satisfaction is computed by means of a fuzzy function (also known as membership function of the fuzzy set defined by predicate P), which is defined for each preference over each item attribute. This section shows how these functions are defined for different types of preferences.

These fuzzy functions have been designed taking into account the considerations listed below. Most considerations constitute design guidelines whose application should result in fuzzy functions modelling intuitive human satisfaction (i.e., buyers' and providers' intuitive evaluations). Additionally, three different examples of the resulting fuzzy functions are presented.

Design guidelines for attribute satisfaction fuzzy functions

- Value preferences for each item attribute define a fuzzy function, whose universe is defined to be the attribute value domain specified in the corresponding item template (offered values do belong to the same domain).
- Satisfaction values belong to the $[0,1]$ interval. When an offered attribute value o_i coincides with a preference, its satisfaction degree is 1. Otherwise, it will take decreasing values (down to 0) as o_i goes further away from the preference inside the domain.

This implies that satisfaction must behave asymptotically when the domain is not limited (that is, for *Any Number* template attribute type).

- If no slope information is associated to a preference (i.e. neither 'LIB -Less is better-' nor 'MIB -More is better-' slopes have been defined), satisfaction must behave symmetrically (and, by following the previous consideration, must assign satisfaction values that decrease proportionally with the distance to the preferred values). For example, if the preferred value is 4, both 3 and 5 offered values should take the same satisfaction degree.
- Multiple values within a preference mean that they are different preferred options (i.e., the preference would be satisfied with any of the values). Thus, all values in a Flat preference interval (or preference set) take the maximum satisfaction:

$\forall o_i \in \text{Preference interval or Preference Set, Satisfaction}(o_i) = 1$

- When preferences are specified through intervals with non-Flat slopes (i.e., MIB and LIB), satisfaction values range from ∞ to 1, being $1 \geq \alpha \geq 0$ the minimum satisfaction degree that preferred values can take inside the preferred interval.

$\forall o_i \in \text{Preference MIB or LIB interval, Satisfaction}(o_i) \geq \alpha$

In our implementation α has been set to 0.5. The value has been fixed because of usability reasons (to relieve users from specifying intricate non-intuitive customisation values).

- Outside non-Flat intervals, those values 'close enough' to the preferred side (i.e., maximum value in an MIB interval or minimum value in an LIB interval) would also take satisfaction values inside $[\alpha, 1]$. By 'close enough' we consider values within a neighbourhood of the interval (computed as a percentage β of the interval length).

Formally, if the preference interval is defined as $[\min_p, \max_p]$ over an attribute whose domain is defined by $[\min_d, \max_d]$ (where \min_d, \max_d are infinite for *Any value* domains), we distinguish two cases:

if preference interval has an MIB slope then:

$$\forall o_i \in (\max_p, \min(\max_d, \max_p + (\max_p - \min_p) \cdot \beta/100)) \text{ Satisfaction}(o_i) \geq \alpha$$

Otherwise, if preference interval has an LIB slope then

$$\forall o_i \in [\max(\min_d, \min_p - (\max_p - \min_p) \cdot \beta/100), \min_p) \text{ Satisfaction}(o_i) \geq \alpha$$

The rest of values always take satisfaction values under α .

$$\text{MIB: } \forall o_i \notin (\max_p, \min(\max_d, \max_p + (\max_p - \min_p) \cdot \beta/100)) \text{ Satisfaction}(o_i) \leq \alpha$$

$$\text{LIB: } \forall o_i \notin [\max(\min_d, \min_p - (\max_p - \min_p) \cdot \beta/100), \min_p) \text{ Satisfaction}(o_i) \leq \alpha$$

Considering neighbourhood intervals is aimed to soften abrupt decreasing (discontinuities) in satisfaction degrees. The value of β been fixed to 10% for usability reasons.

- Simetry does not apply for non-Flat preference intervals. In this manner, satisfaction degrees for offered values falling on the right side outside an MIB interval should decrease with a smoother slope than left-sided values.

The same applies for LIB intervals: left sided values outside the interval should decrease with a smoother slope than right sided values.

- In case an offer contains multiple offered values (o_{i_1}, \dots, o_{i_m}) for an attribute, the satisfaction of the offer for this attribute must correspond to the best offered value satisfaction:

$$\text{Satisfaction}(o_{i_1}, \dots, o_{i_m}) = \max(\text{Satisfaction}(o_{i_1}), \dots, \text{Satisfaction}(o_{i_m}))$$

As previously seen, offers are restricted to take on a single value per attribute. Nevertheless, this consideration has been included so that RFQs can be also scored against provider's preferences using the same scoring computation method. Thus, for example, if a provider's preference for an attribute is 4 and an RFQ requests the attribute values to be 3 or 4, the computed satisfaction must be 1.

Keeping these considerations in mind, we present tree different examples of fuzzy functions defined for different types of attribute domains and preferences (the rest of cases result in analogous fuzzy functions, and thus can be naturally inferred from these ones):

Example 1: Fuzzy function definition for an attribute whose domain is 'range of numbers' and its preference type is 'single'

Figure 7-1 shows the fuzzy function defined for a single preference value of 10 whose template defines its domain as the $[-20, 20]$ range of numbers. The function is generated so that satisfaction is 1 at 10 and it decreases down to 0 at the further extreme of the domain interval

(which in this case is -20). Since simetry must be respected (no slope can be defined for single preferences), the same decreasing slope is applied to values at the right side of the preferred value. In this manner, the maximum value for the domain –wich is 20- will take a satisfaction degree of 0.66.




Figure 7-1: Fuzzy function for a single preference inside a range of numbers domain.

Example 2: Fuzzy function having domain attribute type = ordered set of labels and preference type = flat interval.

This second example corresponds to a fuzzy function whose domain is an ordered set of labels {QA, QB,..., QI} and whose preference is the flat interval [QC, QE] (that is, all QC, QD, QE values are equally preferred) (see Figure 7-2). In this case, labels inside the interval have satisfaction equal to 1, and as in the previous example, satisfaction values decrease for labels down to 0 at QI (the further label). Again, labels on the left of the preferred interval decrease with the same slope than the right side, so QA is assigned a satisfaction degree of 0.5.

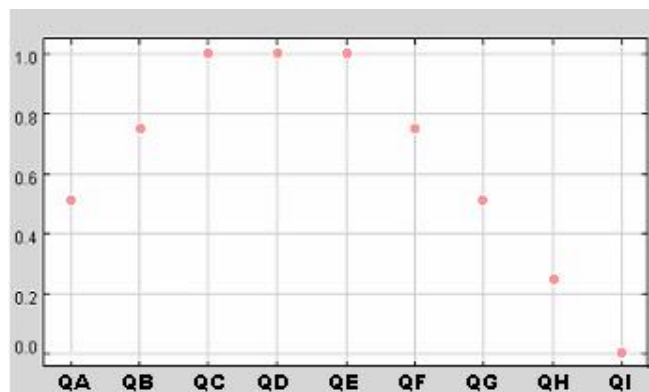


Figure 7-2: Fuzzy function for a flat interval preference and ordered set of labels domain

Example 3: Domain attribute type: any number, preference type: MIB interval.

Last example corresponds to an MIB (more is better) interval preference ([-20,20]) considering a non-restricted numerical domain (that is, any number is allowed). Satisfaction values for this domain are computed using a 4-piece formula considering $\alpha=0.5$ and $\beta=10$ (see Figure 7-3) :

$$(1) \quad \forall o_i \in [-20, 20], \text{ Satisfaction}(o_i) = 0.5 + \frac{o_i + 20}{80}$$

$$(2) \quad \forall o_i \in [20, 24], \text{ Satisfaction}(o_i) = 1 - \frac{o_i - 20}{8}$$

$$(3) \quad \forall o_i > 24, \text{ Satisfaction}(o_i) = e^{-\frac{(o_i - m)^2}{2 \cdot s}} \text{ where:}$$

$$\sigma = \text{spread} \cdot (\text{Length}^2([-20, 20] \cup [20, 24]))^2 \cdot \rho$$

$$\mu = 24, \text{ spread} = 2, \rho = 2$$

$$(4) \quad \forall o_i < -20, \text{ Satisfaction}(o_i) = e^{-\frac{(o_i - m)^2}{2 \cdot s}} \text{ where:}$$

$$\sigma = \text{spread} \cdot (\text{Length}([-20, 20] \cup [20, 24]))^2 \cdot \rho$$

$$\mu = -20, \text{ spread} = 2, \rho = 1$$

First piece (1) corresponds to the satisfaction inside the preference interval, and increases linearly from 0.5 at the minimum up to 1 at the maximum.

Second piece (2) corresponds to the 10% additional interval, which in this example is 4 units long, and decreases from 1 at the maximum of the preference interval (i.e., 20) down to 0.5 at the maximum of the additional 10% interval (i.e., 24).

Third and fourth pieces define satisfaction degrees for the domain values outside the previous intervals. Since the domain is not restricted, we use probabilistic functions that decrease asymptotically along the x axis. Mean (μ) values are set to be the preference interval point that is closer to the described side, (assuming the interval to be the union of both defined preference and additional neighbourhood intervals). In this manner, for the right side (3) mean is the maximum of the additional neighbourhood interval (i.e., 24) and for the left side (4) it is the minimum of the defined preference interval (i.e., -20).

Regarding variance (σ) values, they are assigned as the product of a spread value, a ρ value, and the square value of the length of the interval. Spread is fixed to 2 for both sides, but since symmetry is not desired for an MIB interval, ρ takes different values for each side. More concretely, ρ is set to 2 on the right side (3) and to 1 on the left side (4), so that variance in the side closer to the more preferred values is twice the variance of the other side.

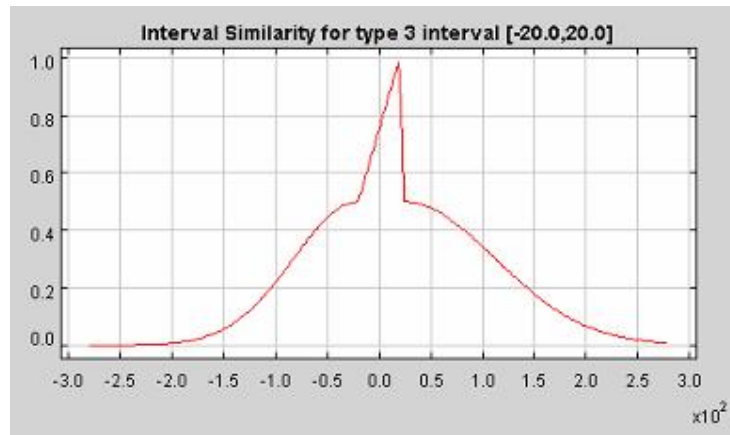


Figure 7-3: fuzzy function for an MIB interval preference and any number domain.

7.1.2 Offer evaluation

An RFQ specifies buyer's preferences over a set of items that describe products (or services) in terms of attributes. Quotes scores providers' offers for an RFQ so that the buyer can have an ordered list of offers, ranking first the offers that best satisfy his preferences. Therefore, in order to score an offer in relation to an RFQ, it is necessary to compute satisfaction values at the level

² $\text{Length}([\min, \max]) = \max - \min$ being $[\min, \max]$ an interval

of attributes and propagate them to the item level and up to the offer level (since an offer may encompass several items).

In a more formal manner, given an offer o composed of m items o^j $j=1..m$, we describe an offer item j as the n -tuple of the form $\langle o_{j_1}^j, o_{j_2}^j, \dots, o_{j_n}^j \rangle$ where $o_{j_i}^j$ represents the i -th attribute value of item j .

In order to compute attribute scoring, Quotes creates the above-described fuzzy functions. Offered attribute values, which are always single, are then used to compute satisfaction degrees:

$$(5) \quad \text{scoring}(o_{j_i}^j) = \text{Satisfaction}(o_{j_i}^j)$$

Attribute satisfaction degrees are then weighted to obtain offer item scoring. Since preferences are RFQ attribute values, weights $w_{j_i}^j$ are taken to be the importance the buyer assigned to an attribute i in the j -th RFQ item (see subsection 2.2).

$$\text{scoring}(o^j) = \sum_{i=1}^{j_n} \frac{\text{scoring}(o_i^j) \cdot w_i^j}{j_n}$$

Finally, we compute offer scoring as a weighted combination of item scorings.

$$s^b(o) = \text{scoring}(o) = \sum_{j=1}^m \frac{\text{scoring}(o^j) \cdot w^j}{m}$$

where w^j stands for the weight of the j -th item.

In the current implementation, all item weights are set to 1 and scoring values are translated from the $[0,1]$ to the $[0,100]$ range.

7.1.3 RFQ evaluation

Similarly to offer scoring functionality, Quotes facilitates providers an RFQ scoring computation that prioritises buyers' requests. Again, fuzzy functions are generated for each attribute, using its domain as defined by templates along with preferences as providers' preferred values (see subsection 3.1).

The main difference appears in the computation of the formula number (5). The reason is that attribute values in RFQ items are not restricted to be single, so that value intervals or sets must be evaluated against these fuzzy functions. In this case, Quotes assumes attribute scoring to be the best satisfaction degree of all values x defined in the attribute value range rfq_i^j .

$$\text{scoring}(rfq_i^j) = \underset{\forall x \in rfq_i^j}{MAX} (\text{Satisfaction}(x))$$

7.2 Winner determination in combinatorial scenarios

Allowing providers to bid on combinations of goods has the interesting feature of enhancing economic/service efficiency. Thus suppliers may offer more competitive bids with the perspective of gaining more business (Rothkopt et al., 1995). However, the determination of an optimal winner combination is a complex problem which, excluding very small instances, cannot be manually solved with common data analysis tools. Thus, consider the decision problem faced by a buyer when negotiating with providers. While in direct auctions, the items that are going to be sold are physically concrete (they do not allow configuration), in a negotiation event involving multiple, highly customisable goods, buyers need to express relations and constraints between attributes of different items. On the other hand, it is common practice to buy different quantities of the very same product from different providers, either for safety reasons or because offer aggregation is needed to cope with high-volume demands. This introduces the need to express business constraints on the number of providers and the amount of business

assigned to each of them. Not forgetting the provider side, providers may also impose constraints or conditions over their offers. Offers may be only valid if certain configurable attributes (e.g. quantity bought, delivery days) fall within some minimum/maximum values, and assembly or packing constraints need to be considered. Once the buyer collects offers, he is faced with the burden of determining the winning offers. The problem is essentially an extension of the combinatorial auction (CA) problem, which can be proven to be NP (Rothkopt 1995). Hence that it would be desirable to relieve buyers from solving such a problem. However, although the application of combinatorial auctions (CA) to e-procurement scenarios (particularly reverse auctions) may be thought as straightforward, the fact is that we identify multiple new elements that need to be taken into consideration.

Current CA reviewed do not model these features with the exception of (Collins et al. 2001) and (Sandholm et al. 2002), where coordination and procurement constraints can be modelled. The rest of work focuses more on computational issues (CA is an NP-hard problem (Rothkopt 1995)) than in practical applications to e-procurement. Suppose that we are willing to buy 200 chairs (any colour/model is fine) for the opening of a new restaurant, and at that aim we employ an e-procurement solution that launches a reverse auction. If we employ a state of the art CA solver, a possible resolution might be to buy 199 chairs from provider A and 1 chair from provider B, simply because it is 0.1% cheaper and it was not possible to specify that in case of buying from more than one supplier a minimum of 20 chairs purchase is required. On the other hand the optimum solution might tell us to buy 50 blue chairs from provider A and 50 pink chairs from provider B. Why, because although we had no preference over the chair's colour, we could not specify that regarding the colour chosen all chairs must be of the same colour. Although simple, this example shows that without means of modelling these natural constraints, solutions obtained are seen as mathematically optimal, but unrealistic and with a lack of common sense, thus obscuring the power of decision support tools, and preventing the adoption of these technologies in actual-world settings.

In this section we present iBundler, a decision support component acting as a combinatorial negotiation solver (solving the winner determination problem) for both multi-item, multi-unit negotiations and auctions. Thus, the component can be employed by both buyers and auctioneers in combinatorial negotiations and combinatorial reverse auctions (Sandholm 2002) respectively. Furthermore, it extends current combinatorial auction models by accommodating both operational constraints and attribute-value constraints.

7.2.1 Desiderata

Next we detail the capabilities required by buyers in the kind of negotiation scenario outlined above. The requirements below are intended to capture buyers' constraints and preferences and outline a powerful bidding language for providers:

1) *Negotiate over multiple items.* A negotiation event is usually started with the preparation of a request for proposal (RFQ) form. The RFQ form describes in detail the requirements (including attribute-values such as volume, quality specifications, dates as well as drawings and technical documentation) for the list of items (goods or services) defined by the negotiation event.

2) *Offer aggregation.* A specific item of the RFQ can be acquired from several providers simultaneously, either because not a single provider can provide with the requested quantity at requested conditions or because buyer's explicit constraints (see below).

3) *Business sharing constraints.* Buyers might be interested to restrict the number of providers that will finally trade for a specific item of the RFQ, either for security or strategical reasons. It is also of usual practice to define the minimum amount of business that a provider may gain per item.

4) *Constraints over single items.* Every single item within an RFQ is described by a list of negotiable attributes. Since: a) there exists a degree of flexibility in specifying each of these attributes (i.e. several values are acceptable) and b) multiple offers referring the very same item can be finally accepted; buyers need to impose constraints over attribute values. An example of this can be the following: suppose that the deadline for the reception of certain item A is two weeks time. However, although items may arrive any day within two weeks, once the first units arrive, the rest of units might be required to arrive in no more than 2 days after.

5) *Constraints over multiple items.* In daily industrial procurement, it is common that accepting certain configuration for one item affects the configuration of a different item, for example, when dealing with product compatibilities. Also, buyers need to express constraints and relationship between attributes of different items of the RFQ.

6) *Specification of providers' capacities.* Buyers cannot risk to award contracts to providers whose production/servicing capabilities prevent them to deliver overcommitted offers. At this aim, they must require to have providers' capacities per item declared.

Analogously, next we detail the expressiveness of the bidding language required by providers. The features of the language below are intended to capture providing agents' constraints and preferences.

7) *Multiple bids over each item.* Providers might be interested in offering alternate conditions/configurations for a same good, i.e., offering alternatives for a same request. A common situation is to offer volume-based discounts. This means that a provider submits several offers and each offer only applies for a minimum (maximum) number of units.

8) *Combinatorial offers.* Economy efficiency is enhanced if providers are allowed to offer (bid on) combination of goods. They might lower the price, or improve service assets if they achieve to get more business.

9) *Multi-unit offering.* Each provider needs to specify that they will only participate in trading if a minimum (maximum) amount of business is assigned to him.

10) *Homogeneous combinatorial offers.* Combinatorial offering may produce inefficiencies when combined with multi-unit offering. Thus a provider may wind up with an award of a small number of units for a certain item, and a large number of units for a different item, being both part of the very same offer (e.g. 10 chairs and 200 tables). It is desirable for providers to be able to specify homogeneity with respect to the number of units for complementary items.

11) *Packing constraints.* Packing units are also a constraint, in the sense that it is not possible to serve an arbitrary number of units (e.g. a provider cannot sell 27 units to a buyer because his items come in 25-unit packages). Thus providers require to be capable of specifying the size of packing units.

12) *Complementary and exclusive offers.* Providers usually submit XOR bids, i.e., exclusive offers that cannot be simultaneously accepted. Also, there may exist the need that an offer is selected only if another offer is also selected. We refer to this situation as an AND bid. This type of bids allows to express volume-based discounts. For example, when pricing is expressed as a combination of base price and volume-based price (e.g. first 1000 units at €2.5 per unit and then €2 each).

Obviously, many more constraints regarding pricing and quantity can be considered here. But we believe these faithfully address the nature of the problem. Actually, iBundler has been applied to scenarios where some of these constraints do not apply while additional constraints needed to be considered. This was the case of a virtual shopping assistant, an agent that was able to aggregate several on-line supermarkets and optimise the shopping basket. To do so, it was necessary to model the fact that delivery cost depends on the amount of money spent at each supermarket.

7.2.2 Implementing winner determination

Consider the problem faced by a buyer aiming at choosing the optimal set of offers sent over by providers taking into account the features of the negotiation scenario described above. The problem is essentially an extension of the combinatorial auction (CA) problem in the sense that it implements a larger number of constraints and supports richer bidding models. The CA problem is known to be NP-complete, and consequently solving methods are of crucial importance. Many of the works reviewed in the literature adopt global optimal algorithms as a solution to the CA because of the drawbacks pointed out for incomplete methods. Basically two approaches have been followed: traditional Operations Research (OR) algorithms and new problem specific algorithms (Fujishima et al. 1999 and Sandholm 2002). It is always an interesting exercise to study the nature of the problem in order to develop problem specific algorithms that exploit problem features to achieve effective search reduction. However, the fact is that the CA problem is an instance of the multi-dimensional knapsack problem MDKP (as

indicated in (Holte 2001)), a mixed integer program well studied by the operation research literature. It is not surprising that, as reported in (Anderson et al. 2000), many of the main features of these problem specific new algorithms are rediscoveries of traditional methods in the operations research community. In fact, our formulation of the problem can be regarded as similar to the binary multi-unit combinatorial reverse auction winner determination problem in (Sandholm et al. 2002) with side constraints (Sandholm et al. 2001). Besides, expressing the problem as a mixed integer programming problem with side constraints enables its resolution by standard algorithms and commercially available, thoroughly debugged and optimised software which have shown to perform satisfactorily for large instances of the CA problem.

With these considerations in mind, the core of the iBundler decision support service has been modelled and implemented as a mixed integer programming (MIP) problem. We have implemented two versions: one using ILOG CPLEX 7.1 in combination with SOLVER 5.2; and another one using iSOCO's Java MIP modeller that integrates the GLPK library (Makhorin 2001). In both cases it takes the shape of a software component.

7.2.3 Winner determination performance

Since combinatorial auction solvers are computationally intensive, a major issue is whether our service is to behave satisfactorily in highly-demanding trading scenarios. At this aim, we have conducted some empirical measures on the performance of iBundler. Figure 7-4 shows how *it* behaves when solving negotiation problems as the number of bids, the number of items, and the items per bid increase. Notice that in order to run our tests, we devised a customisable generator of data sets which artificially created negotiation problems by wrapping the solution with noisy bids. In this way, not only were we able to measure the performance of iBundler, but also to automatically verify the sound behaviour of the service in a large variety of negotiation scenarios demanding the many capabilities of the service.

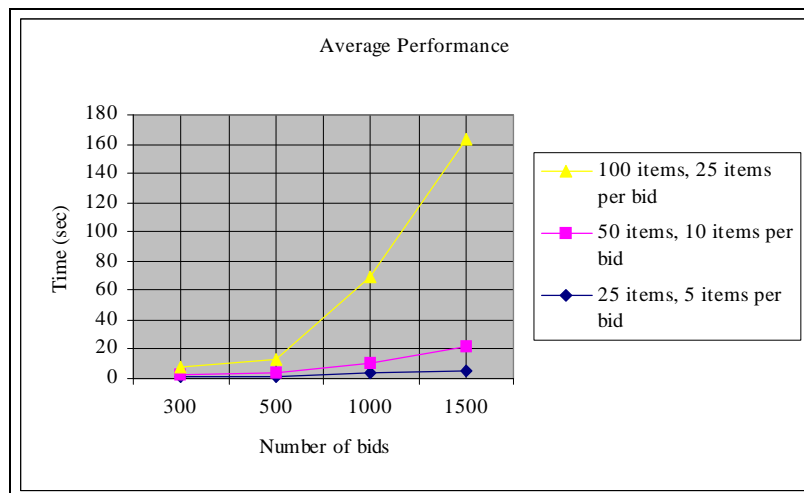


Figure 7-4. Average performance of the iBundler decision support component

8 Multi-stage sourcing events

Once the different tasks that are usually involved in sourcing events have been detailed in previous sections, this section aims to describe Quotes support for multi-stage negotiations. (E-negotiation group 2000)

A typical situation in industry is the specification of sourcing events that involve several stages. Depending on whether the product features are not fully known by the buyer, the product's complexity and/or the number of potential providers, sourcing professionals design the sourcing process as a sequence of steps (stages) with various objectives.

Consider the negotiation of a new frame contract for raw material:

- Typically the sourcing event will start with new providers being invited to an RFI (request for information) with the goal to qualify them. An RFI negotiation is performed analogously to RFQ negotiations (it requires an RFI elaboration task similar to RFQ elaboration as well as providers' responses). The result of this stage is the determination of winning (qualified) providers. Notice that already qualified providers will not be invited to this stage.
- At the end of the RFI stage, qualified providers are then invited to an RFQ with the aim of negotiating required product features and transaction terms. On the providers' side, offers are either manually or automatically built as responses to the received RFQ by assigning values to the product features and transaction terms. Thereafter, the buyer can conduct simultaneous one-to-one negotiations as part of the one-to-many negotiation process.
- This negotiation phase may end up with the buyer accepting some offer(s) or may be continued as a reverse auction stage with a subset of the most competitive providers. At this point, the buyer may opt for launching a reverse auction with the objective of increasing competitiveness along those product features and transaction terms whose negotiated values remain unsatisfactory.

As we have observed in current market solutions there is little support for creating customized sourcing events as compositions of negotiation processes. Quotes allows buyers to undertake multi-stage sourcing events by means of stage switching as described above. Moreover RFI, RFQ, and auction stages can be interleaved (and repeated) by the buyer at will. Conditions and product requirements can be changed from stage to stage, for instance, attribute importance adjustments or new product requirements. Offers can be migrated from stage to stage as long as they still satisfy these new requirements. Additionally, this will help the buyer to determine providers that should be excluded out of the next stage (although in some cases providers may still be invited but only to compete for a subset of a multi-item RFQ). Obviously, new providers can always be invited to a next stage.

Quotes controls the status of the whole process as well as the state of each provider, thus communicating to them their inclusion or exclusion in the next stage.

Figure 8-1 shows how Quotes implements multi-stage sourcing events.

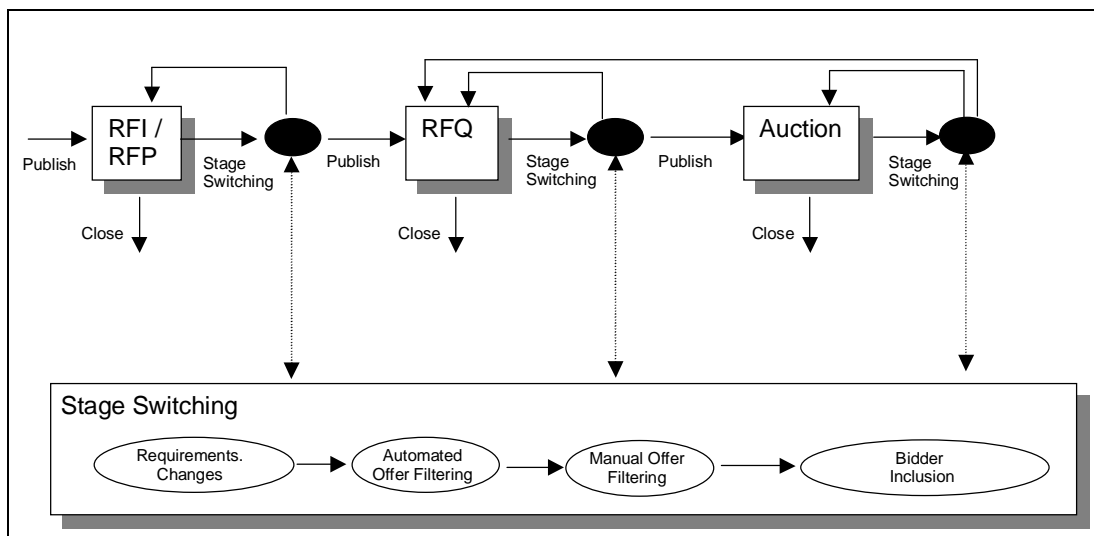


Figure 8-1. Multi-stage sourcing at Quotes

A sourcing event starts with the definition of a new RFx (RFI, RFP, RFQ and reverse auction mechanism). The RFx is published and executed. When the RFx finishes the buyer can determine winners and close the stage, thus finalizing the sourcing event. Alternatively the

buyer can undergo a *switch-to-stage* procedure. By doing so, Quotes creates a new RFX stage based on the current one. The buyer can then change the requirements (add or remove items, add or remove attributes) and determine the negotiation mechanism (RFI, RFP, RFQ or auction). Quotes automatically filters out (rejects) offers that do not satisfy newly requirements and allows the buyer to manually cancel others. Finally, the buyer decides which providers continue with the next event and has the choice of inviting new ones.

9 Results

Quotes has been successfully applied to several real-life negotiation scenarios of varying complexity (from single line to multi-line, multi-attribute) and economic value (few thousand € to frame contracts of several million €). Next we summarize the most remarkable outcomes.

- Actual buyers encountered Quotes scoring mechanism as a simple, intuitive, and powerful way to quickly differentiate good from bad offers. However, when making final decisions or comparing similar offers they often obviated the scoring values provided by Quotes and fine-analysed offers by means of other evaluation functions. We are currently trying to incorporate alternative scoring functions that cope with the full needs of sourcing professionals along with preference elicitation mechanisms (although this issue remains the Achilles tendon of sourcing applications).
- A common agreement was the convenience of modelling off-line negotiation processes in a natural way, without introducing inefficiencies and frictions derived from changing the “rules of the game” (that is, for example, substituting the off-line negotiation processes by on-line auctions or using negotiation artefacts that do not model previous processes). Furthermore, providers appreciated the transparency introduced by the tool (since all participants actions can be audited).
- Buyers belief is that combinatorial offering introduces high complexity for providers to bid and to understand auction dynamics. Consequently, Quotes’ combinatorial capabilities have been solely applied to small set of actual-world scenarios where the buyer pre-defines valid item combinations on which providers can bid.
- The possibility of automatic offer submission is seen with interest for repetitive sourcing events in private e-sourcing platforms where providers and business rules are well known or belong to a provider qualification procedure or a frame contract. Nonetheless, the full application of such automatisms faces cultural barriers such as providers being not so keen on revealing capabilities/preferences to third parties; perception of e-sourcing tools as a hazard for sourcing professionals, etc.
- Leading users suggested additional auction rules that best suited their necessities. For example, a buyer forced Quotes to incorporate a bidding rule that resulted in an increase of the number of participants in an auction event. This rule allowed inactive bidders to send their first bid without overbidding the best bid.

Finally, the results obtained in terms of economic outcome were no different than the promises made by e-sourcing analysts. Negotiation time was reduced from weeks to days, mostly due to the elimination of communication synchronism (telephone, fax) and administrative tasks. Price and condition benefits were also obtained. Obviously, price savings were more noticeable in auctions (13,6% in average), but on-line negotiation also achieved price/service reduction below target, a result that increased buyer’s satisfaction with the tool.

10 Conclusions

E-sourcing is becoming an established part of the business landscape (Stephens Inc., 2001). We are witnessing the continuous, tightly competitive progress of e-sourcing applications in the market. Nonetheless we can still identify two major, unsatisfactorily solved issues that prevent them from supporting effective strategic sourcing, namely:

- *Capability to support sourcing processes for varying industries and businesses.* Since sourcing processes are highly dependent on each business case (because of each industry's particularities and individual businesses' practices) it is extremely complex to capture in a single product all processes and negotiation requirements of a general-purpose solution. Along this direction, perhaps the major drawback of most market solutions is the lack of support for creating customised sourcing events as compositions of negotiation processes (multi-stage negotiation processes). For instance, a sourcing event might be composed of a pre-defined sequence of auctions or an interleave of auctions and negotiations. In this work we have presented an e-sourcing solution that does indeed support customisable multi-stage negotiations.
- *Decision support tools for strategic thinking.* Strategic sourcing is founded on the availability of powerful decision support tools. Nonetheless current vendors' solutions are lacking as to this matter. As an example, most tools share the commonality of not providing support for determining the winner in multi-item, multi-attribute negotiations and auctions. The unavailability of such support poses an intricate, combinatorial problem to professional buyers that leads them to either relinquish or opt for alternative, and less efficient, non-combinatorial protocols. In this work we have tried to make headway in providing decision support

In this paper we have tried to exemplify how the sourcing process can be highly automated, allowing companies to achieve enormous benefits: cost savings, processing time reduction, less time-to-market, and more time left to strategy. We have presented our contribution along this direction by dissecting Quotes, an Internet-enabled sourcing solution capable of streamlining the sourcing process. Quotes' main strengths can be summarised as follows:

- It allows goods and services to be represented and managed with all their attributes, overcoming rigid and unreal price-discovering approaches.
- It provides a powerful negotiation framework based on the composition of structured negotiation protocols and flexible reverse auctions.
- It provides with the necessary tools to help users manage the complex sourcing mechanisms involved in multi-item, multi-unit, multi-attribute, multi-stage negotiations.

The fundamental contribution of Quotes lies on the incorporation of highly valuable decision-making support functionalities targeted at spurring the transition of sourcing processes from transactional to strategy-centered. This paper describes three decision support systems that have been studied and implemented in Quotes with satisfactory results, namely:

- *Automated offer generation.* Providers can translate their business knowledge into *bidding* rules that allow instantaneous and automatic construction of recommended offers.
- *Scoring algorithm.* Based on the importance assigned by the buyer to each item attribute in an RFQ and his flexibility to accept offers beyond his preferences, a *fuzzy* offer-matching algorithm scores each offer and ranks it accordingly. Analogously, providers can benefit from the very same algorithm in order to discriminate incoming RFQs.
- *Computation of the optimal bid set for combinatorial negotiations.* Given a set of offers for either a multi-item RFQ or an auction, the iBundler component obtains the optimal subsets of offers according to various criteria (f.i. minimization of price, start/finish date, maximization of product quality, etc.), or a combination of them specified through scoring. Although the used algorithm is not new, the way the problem is formulated can be considered as an important contribution to research in combinatorial negotiations.

Finally, there is an important aspect affecting the successful exploitation of e-sourcing applications: negotiation design matters. Indeed there are many dimensions involved in the design of negotiation scenarios. Thus we must keep in mind that careless, faulty designs of such scenarios may eventually lead to terribly catastrophic outcomes as reported through several case studies in (McAfee, 1995).

References

- [1] Fujishima, Y.; Leyton-Brown, K.; Shoham, Y. (1999) Taming the computational complexity of combinatorial auctions: Optimal and approximate approaches. In *Proceeding of the Sixteenth International Joint Conference on Artificial Intelligence, IJCAI'99*, 548–553.
- [2] Keeny, R. L.; Raiffa, H. (1993). *Decision Making with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge, UK: Cambridge University Press.
- [3] Korf, R. (1998). Artificial intelligence search algorithms, *CRC Handbook of Algorithms and Theory of Computation*, M.J. Atallah (Ed.), CRC Press, Boca Raton, FL, pp. 36-1 to 36-20.
- [4] Ribeiro, R. A. (1996). Fuzzy multiple attribute decision making: A review and new preference elicitation techniques. *Fuzzy Sets and Systems*, 78; 155-181.
- [5] Rothkopt, M. H.; Pekec, A.; Harstad, R. M. (1995). Computationally manageable combinatorial auctions. *Management Science*, 44(8):1131-1147.
- [6] Sandholm, T. W. (2002) "Algorithm for optimal winner determination in combinatorial auctions". *Artificial Intelligence Journal* 135:1-54, 2002.
- [7] Sandholm, T. (2000). eMediator: A Next Generation Electronic Commerce Server. International Conference on Autonomous Agents (AGENTS), Barcelona, Spain, June 3-8.
- [8] T. Sandholm, S. Suri, A. Gilpin, and D. Levine (2002). "Winner Determination in Combinatorial Auction Generalizations". First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'02), Bologna, Italy, 2002, pp. 69-76.
- [9] De Vries, Sven; Vohra, Rakesh. (2001). Combinatorial Auctions: A survey.
- [10] Godo, L.; López de Mántaras, R.(1993). Fuzzy Logic. In *Encyclopedia of Computer Science and Technology*. Marcel Dekker, Inc. Pennsylvania. Vol 29, pp 211-229.
- [11] Ariba. (2002). Ariba Enterprise Sourcing feature description available on the web at <http://www.ariba.com>.
- [12] Commerce One. (2002). Commerce One Source features description available on the web at <http://www.commerceone.com>.
- [13] MetalSite (2003): the auto-referenced Internet's premier site for buying and selling metals at <http://www.metalsite.com/>.
- [14] Reyes-Moro, A.; Rodríguez-Aguilar, Juan A.; Cerquides, Jesus; López-Sánchez, Maite; Gutiérrez-Magallanes, David. (2002). Quotes: a Negotiation Tool for Industrial Procurement. Fourth International Conference on Enterprise Information Systems (ICEIS 2002), April 4-6, Ciudad Real, Spain.
- [15] A. Reyes-Moro, J.A. Rodríguez-Aguilar, M. López-Sánchez, J. Cerquides, and D. Gutierrez-Magallanes (2003). "Embedding decision support in e-sourcing tools: Quotes, a case study". *Group Decision and Negotiation Journal*, 2003 (to appear).
- [16] A. Reyes-Moro, J.A. Rodríguez-Aguilar, J. Cerquides, M. López-Sánchez, and D. Gutierrez-Magallanes (2001). "Quotes, the way to the best deal" Invited talk at The DEXA 2001 Workshop on e-negotiations, 5-8 September 2001, Munich.
- [17] Aberdeen Group. (2002). Making E-Sourcing Strategic: From Tactical Technology to Core Business Strategy.
- [18] McAfee, R. Preston. (1995). Auction Design for the Real World. <http://www.eco.utexas.edu/faculty/McAfee/Papers/PDF/Auction.pdf>.
- [19] Stephens Inc. (2001). Internet Supply Chain Team. Strategic Sourcing: Applications to Turn Direct Materials Procurement into Competitive Advantage. Industry Report, January 30.
- [20] Kelly, F.; Steinberg, R.(2000). [A combinatorial auction with multiple winners for universal service](#). *Management Science* 46:586-596.
- [21] Zurel, E.; Nissam, N. (2001). An Efficient Approximate Allocation Algorithm for Combinatorial Auctions. Proceedings of the ACM Conference on Electronic Commerce, Tampa, Florida.

- [22] Hoos, H. H.; Boutilier, C. (2000). [Solving Combinatorial Auctions using Stochastic Local Search](#). In *Proceedings of AAAI-2000*.
- [23] Holte, R. C. (2001). *Combinatorial Auctions, Knapsack Problems, and Hill-Climbing Search*. Lecture Notes in Computer Science 2056.
- [24] Andersson, A.; Tenhunen, M.; Ygge, F. (2000). *Integer programming for combinatorial auction winner determination*. In *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS)*, pages 39--46, Boston, MA.
- [25] Collins, J.; Gini, M. (2001). An Integer Programming Formulation of the Bid Evaluation Problem for Coordinated Tasks. In B. Dietrich and R. V. Vohra, editors, *Mathematics of the Internet: E-Auction and Markets*, volume 127 of IMA Volumes in Mathematics and its Applications, pages 59-74. Springer-Verlag, New York.
- [26] Bichler, M.; Kalagnanam, J.; Lee, H. S. (2002): RECO: Representation and evaluation of configurable offers. In: Hemant K. Bhargava and Nong Ye (Eds.): *Computational Modeling and Problem Solving in the Networked World: Interfaces in Computing and Optimization*, Kluwer Academic Publishers, 2002.
- [27] Bartel Van de Walle, Sven Heitsch and P. Faratin (2001): Coping with One-to-Many Multi-Criteria Negotiations in Electronic Markets: *in Proceedings of the Second International Workshop on Negotiations in Electronic Markets: beyond Price Discovery, held during DEXA 2001* (Munich, Germany, September 2001), 747--751.
- [28] Baas and Kwakernaak (1977), Rating and ranking of multiple-aspect alternatives using fuzzy sets, *Automatica* 13 (1977) 47-58.
- [29] Makhorin, A.(2001) "GLPK – GNU Linear Programming Toolkit". <http://www.gnu.org/directory/GNU/glpk.html>. 2001
- [30] E-negotiation group (2000) "The London Classification" at DEXA workshop on e-negotiations. London, 2000.
- [31] Kalagnanam, J. and Parkes, D. C. (2003) "[Auctions, Bidding and Exchange Design](#)". Chapter 10, in *Supply Chain Analysis in the eBusiness Era*, Edited by David Simchi-Levi, S. David Wu and Z. Max Shen, 2003. (forthcoming)