

Towards Agent-Based Service Composition Through Negotiation in Multiple Auctions

Christ Preist, Andrew Bye, Claudio Bartolini and Giacomo Piccinelli

Hewlett-Packard Laboratories
Filton Road, Stoke Gifford, Bristol, BS32 8QZ, United Kingdom
*Chris.Preist@hp.com ; Andrew.Byde@hp.com ;
Claudio.Bartolini@hp.com ; Giacomo.Piccinelli@hp.com*

Abstract

Service composition is the act of taking several component products or services, and bundling them together to meet the needs of a given customer. In the future, service composition will play an increasingly important role in e-commerce, and automation will be desirable to improve speed and efficiency of customer response. In this paper, we discuss the technical issues surrounding the automation of dynamic electronic service composition, using a fictitious company, FreightMixer, to demonstrate the process. We focus specifically on the issue of appropriate negotiation strategies for service composition, and present the specification of an algorithm to provide a robust solution to these problems in the context of multiple simultaneous auctions. We present a worked example to demonstrate the behaviour of the algorithm, and discuss related and future work.

1 Introduction

Over the past few years, Electronic Commerce has become an increasingly central part of the economy. An Internet presence is considered an essential part of doing business, rather than an exotic add-on to a company. More and more transactions, both from business to consumer and between businesses, are taking place online. Simple fixed cost business transactions are often automated at one or both ends, and auctions are overwhelmingly conducted by automated auctioneer software. Agent technology has been proposed as a means of automating some of the more sophisticated negotiations which businesses are involved in (e.g. (Jennings et al., 1996)). In this paper we look at a specific class of business process that will become increasingly important in the virtual economy - service composition. We consider the different technical issues that must be addressed if service composition is to be automated, and focus specifically on algorithms for the purchase of composite services from a group of auctions.

Over the last decade, companies have been encouraged by business consultants (Peters et al., 1984) to focus on their core competences. By trying to do everything - welding, graphic design, supply chain management, customer care, keeping the photocopiers running, producing good food in the office canteen - companies run the risk of being 'jack of all trades, but master of none'. As a result of this there is a danger that other smaller companies focused on the same core business will outperform them. To avoid that risk, and become more competitive, large companies are going through a process of 'disaggregation'. In some cases, this can mean splitting a large company into several parts, each

of which can focus on one core business (such as the recent move by Hewlett Packard to separate its test and measurement business from its computing business, creating a new company, Agilent, from the former.) In other cases, it can mean outsourcing more and more of a companies activities to other companies, maintaining only those activities that it truly excels in.

This trend is beginning to have an impact on many E-businesses, as well as traditional bricks-and-mortar companies. Companies would like to be able to outsource some of their activities over the Internet. Initially, this has focused on semi-permanent arrangements, with the web acting as an intermediary. (For example, career guidance information is provided to HP employees via a web-based third party.) However, as this trend is becoming increasingly important, much research and development effort has been focusing on a new vision for the Internet - e-services. E-services are virtual entities that provide a service over the network through an open standard interface. The service may be information, such as the latest stock prices, or may be a virtual representation of some physical good or activity, such as a contract to transport a crate from one location to another. Because the service is offered through an open standard interface, any client familiar with this standard can use it. Furthermore, the output from one service can be fed directly into another service. This makes the creation of composite services and complex business processes which cross organizational boundaries possible. Potentially, this can be done automatically and dynamically, and agent technology will play a key role in this.

This leads to the emergence of an important role in the virtual economy - the service composer. As companies focus on their core competencies, other companies can focus on creating composite packages. This is not new - travel agents, among others, have done exactly that for years - but what is new is that it will be able to take place dynamically, automatically, over the Internet. In this paper, we discuss the technical issues that must be overcome if this is to come about, and focus specifically on negotiation algorithms. Firstly, in §2, we introduce the problem of service composition, and discuss which technical issues must be overcome if it is to be automated in e-commerce. In §3, we present an example service composition scenario involving a virtual company, "FreightMixer". In §4, we focus specifically on the problem of participating in multiple auctions to purchase service bundles. We present an algorithm specification, and give an example of the algorithms behaviour. We then discuss related work, and finish by presenting conclusions and future work.

2 Issues in Service Composition

In an automated B2B transaction, the participants must go through three conceptually separate phases; matchmaking, negotiation and service delivery. (This lifecycle is an abstraction of that used in (Jennings et al., 1996).) We briefly describe these three phases, and then discuss how an enterprise involved in service composition participates in them. We conclude the section with an example scenario, taken from the freight services domain, that will illustrate the concepts discussed.

2.1 Matchmaking phase

Matchmaking is the process of putting service providers and service consumers in contact with each other. For matchmaking to take place, services that wish to be dynamically located must publish details of themselves, and entities wishing to locate such services

must search for these details. Some of the services advertising themselves for match-making will be simple end-providers that can be negotiated with directly. Others may be brokers, auction houses and marketplaces that offer a locale for negotiating with and selecting among many potential providers offering similar services.

The services advertise a service description; a formal specification of the nature of the service they offer. Usually this information will be held in a central matchmaking directory. The facilitator agents of KQML (Finin and Fritzson, 1994) provide one approach to handling this.

When an entity wishes to locate a service of a certain type, it queries the matchmaker with a service request. This request takes a similar form to the service description, but may have certain fields unbound or constrained. The matchmaker returns a set of pointers to negotiations with appropriate service providers; in some cases, these negotiations are 1-1, while others may be auctions, exchanges, etc. Each pointer may also have a contract template associated with it, showing the associated terms and conditions of the negotiation. (Of course, some of these may be uninstantiated or semi-instantiated, and therefore open to negotiation.) Standardization is essential to allow effective matchmaking. FIPA (Dale and Mamdani, 2001) is currently developing one approach. UDDI provides a less rich, but more widely supported, alternative.

2.2 Negotiation phase

After matchmaking, a service consumer is faced with a variety of potential negotiations. Its aim is to procure the best service, taking into account factors such as price, speed of delivery, etc. To do this, it will participate in one or more of these negotiations. Different negotiations will have different market mechanisms - sets of rules determining how the negotiation should take place. The simplest such mechanism consists of a single service, offering itself at a fixed, non-negotiable, price. Other services may be willing to enter in to 1-1 bargaining with potential customers, or conduct auctions. Others may post their availability through exchanges, together with many other similar services. As a result of this negotiation, the different parties will agree a contract with terms and conditions that give each member certain rights (such as the right to use a certain service) and obligations (such as the obligation to pay a certain price) (Dignum and Weigand, 1995; Tan and Thoen, 1999; Norman and Reed, 2000).

2.3 Service delivery phase

Once the terms and conditions of service execution have been agreed by the participant to negotiation, service execution can start. That involve interaction between the service provider and the service consumer, to act according to the terms and conditions established during the negotiation process.

2.4 Service composition

Service composition is the act of purchasing several *component services*, combining them, and selling them as a single composite service. The *service composer* responsible for the generation of the composite service must purchase the component services from a group of *suppliers* and will sell the composite service to one or more *customers*. In §3, we will give a detailed example of a company responsible for shipping freight. The company, FreightMixer, is the service composer. It is approached by a customer with a request to

ship a crate from London to San Francisco. However, it does not own any cargo facilities of its own. Instead, it subcontracts, and arranges cargo space on a set of linked flights from London to San Francisco. The airlines running these flights are the suppliers, and the individual flights are the component services.

Of course the concepts of component and composite services are relative, based on the perspective of the service composer. A supplier may in turn be a service composer, and view the component service as a composite service from their perspective. Similarly, the buyer may be a service composer, using this service as a component in a larger composite service. For example, the shipping of the crate may be on behalf of a conference venue organiser, who is using the display materials in the crate to prepare a conference in San Francisco. The shipment is a component service, and the conference is the composite service.

The composite nature of the service affects the behaviour of the composer in all phases of the business transaction, and requires some modification to the standard, static view of these phases.

Matchmaking is traditionally viewed as a lookup process to find service providers able to meet a requesters needs, prior to the requester selecting and/or negotiating with them. Service providers simply advertise information about the service they offer in a database, and requesters use this database to make their selection. However, if the virtual economy is to encourage dynamic service composition, more flexibility will be necessary at this stage. A provider will have some idea of the general services it is interested in offering, but will not know the full details. At any given time, it can estimate these details based on the current state of markets. Hence, if it is to participate in matchmaking, the matchmaker must play a more active role. It must route potential service requests to service composers, which then respond with a dynamic service advertisement detailing the closest service to the request they can offer. This advertisement should not be treated as binding - it is simply an estimate based on the current market situation. Negotiation would be required to reach a binding contract.

In a context of dynamic service composition, additional requirements will be imposed on the service description to be advertised by the matchmaker and on the queries the system can deal with. The description should include abstract roles such as “insurance provider”. The customer will know that these will be filled by subcontractors found by the service composer, but it will not know a priori who will take these roles. The service composer will dynamically negotiate with potential subcontractors to determine exactly who will perform these tasks. Service composers must take into account any restrictions that potential customers may wish to place constraints over who plays the roles. For example a customer may want to ensure that all subcontractors are members of appropriate trade bodies. For this reason, the service provider may need to give information to the customers during the matchmaking process about the names and/or details of potential subcontractors.

When the composer is either advertising a composite service or responding to a request through the matchmaker, then it needs models of how to decompose a service description into base service types that it can try to procure. Initially, this will be done at the service specification level. From a declarative description of the high level service, it will generate declarative descriptions of the sub-services which can be used during matchmaking to locate potential subcontractors.

The way in which a service request can be decomposed is not unique. The composer may generate many alternatives in advance, and then place requests for the base services.

Alternatively, it can use the currently advertised services to inform the generation process. Which of these two strategies is appropriate will depend on whether the base services work in a “push” or “pull” advertising mode.

In the negotiation phase, the service composer will be involved in many interlinked negotiations. For any single bundle of base service types, the composer will be involved in at least one, but more likely many negotiations to acquire instances of each service type. Furthermore, the service composer may simultaneously negotiate to purchase alternative bundles, in an effort to find which bundle is best. Ideally, it would like to do this in a non-committing way. However, some forms of negotiations (such as auctions) require participants to make a commitment when placing a bid, and provide no guarantees of success. When participating in negotiations of this kind, the provider of the composite service must take care to avoid buying incomplete or overly large bundles. Furthermore, the service composer may be simultaneously negotiating with potential clients. In this case, it must trade off its expectation of winning such negotiations against any commitments it makes in the negotiations to purchase base services.

As we have seen in the matchmaking phase, the composite service provider relies on the declarative description of the sub-services when generating potential bundles of base services. During negotiation the service provider must ensure that the base services truly can inter operate to provide the composite services. The declarative representation of the services can only guarantee this if the community as a whole defines standards of inter-operability. If this is not the case, the declarative description will have to be refined. In addition to the specification of the services, the parties will need to agree on the protocol that they will use to communicate during service delivery (e.g. what exchange of messages will take place to make a payment) and on the implementation of the interactions between the parties (e.g. what is the format of the messages that are exchanged). For more details on service composition, see (Piccinelli and Lynden, 2000; Piccinelli and Mokrushin, 2001)

At the end of the negotiation phase, the service agreement forms the basis of a contract. Usually contracts will be between two parties. There will be one contract between the requester and the service provider and one contract between the service provider and each subcontractor. However, in some circumstances, multi-party contracts may be appropriate. They provide additional security to the service provider by offloading risk onto the subcontractors and onto the client.

After the contract has been formed, service delivery can commence. During service delivery, it is the responsibility of the service provider that the execution proceeds smoothly. Therefore the service provider will orchestrate the execution flow and ensure that each component service inter-operates appropriately. The orchestration relies upon a monitoring infrastructure that makes use of all the levels of service descriptions - declarative specification, procedural protocols and implementation of interactions. The subcontractors will play roles that appear in these descriptions. In order to fulfill these roles, each subcontractor will obtain an appropriate view over the original description which it will execute.

3 Example scenario: FreightMixer

In order to illustrate these concepts, we now present a scenario to show the issues involved in service composition, and the impact of combined negotiation techniques. The scenario

is taken from the freight domain. FreightMixer is an imaginary transport company that exploits cheap last-minute sales of excess hold space. While it may not be the quickest service, it aims to be the cheapest. Electronic marketplaces are both a source of resources (individual flight legs) and a channel for products (composite flights for a given customer).

FreightMixer does not own any transport infrastructure. Instead, it aims to dynamically design a cost-effective solution, using whatever third-party services are currently available cheaply. It composes these individual services together into a value-added solution which can be offered to customers at a premium. Its business model revolves around the dynamic acquisition of transport services at a competitive price, and the profitable sell of the composite service. Hence, effective negotiation techniques are crucial to the procurement as well as to the sales function of the company.

The knowledge that FreightMixer has of the freight market is the main asset of the company, and very basis on which its competitive advantage is built. Crucial aspects of this knowledge are captured electronically, allowing algorithms to automatically design and implement end-to-end solutions. In particular, it must have domain knowledge about when two flight legs can be linked together, and how to do this. It must know how much time is required to get the crate from one plane to the other, how to contract with appropriate ground staff in different airports to arrange the hand-over, and what paperwork is necessary to enter different airports.

We now apply the three-stage model of business transactions to a typical deal generated by FreightMixer, and discuss what functionality a service composition company requires.

3.1 Matchmaking Phase

During the matchmaking phase, FreightMixer acts in two distinct sets of markets:

In the markets for end-to-end cargo services, it acts as a potential seller. It observes the advertised requirements of potential customers in this market.

In the markets for hold space on flights (and possibly ships), it acts as a potential buyer. It observes the availability and cost of different options in these markets.

In its role as service composer, it must (a) understand requirements of the potential customers which are currently requesting services in the end-to-end cargo markets, and identify a service which could meet their needs (b) identify the alternative ways this service can be created from component services (ie hold space on specific flights) (c) identify potential sellers of these component services in the markets for hold space on flights.

As an outcome of the matchmaking phase, Freightmixer will have a list of negotiation options. Each option will consist of the following:

- A potential buyer, or set of buyers, who are currently requesting a service in the end-to-end cargo marketplaces.
- A service specification which meets the needs of these buyers.
- One or more alternate decompositions of this service into component services.
- A list of sellers in the markets for hold space who are offering to sell individual component services appearing in these decompositions.

For example, assume FreightMixer observes a Request For Quotes reverse auction for sending a 1 tonne crate from London to San Francisco, with the best offer currently

at 210. Using its database of service models, Freightmixer identifies alternative combinations of flights which might potentially meet these needs. It identifies a direct route from London(LHR) to San Francisco(SFO), and also identifies alternative routes via Chicago(ORD), New York(JFK) and Boston(BOS). It then checks the auctions for excess hold space and finds that appropriate auctions exist for all legs except LHR to JFK. The remaining alternatives it has are shown in Figure 1.




Figure 1: Graph of services

Hence, FreightMixer has an option consisting of the buyer conducting the reverse auction, three alternative ways of generating the required service from component services ($\{LHR-SFO\}$, $\{LHR-ORD \& ORD-SFO\}$ or $\{LHR-BOS \& BOS-SFO\}$) and potential sellers for each of the component services. The options would also include subsidiary services (e.g. insurance, re-packaging, temporary storage) which we shall not discuss.

FreightMixer may also choose to pro-actively advertise certain composite services during the matchmaking phase. If it expects demand for certain services, it can go out and provisionally negotiate for the individual components while waiting for clients to respond to the advertisement.

3.2 Negotiation Phase

During the negotiation phase, Freightmixer must again participate in two sets of markets. It must participate as a seller in the markets for end-to-end cargo services, negotiating over the terms and conditions of sale with the various buyers identified in its set of options. It must participate as a buyer in the markets for hold space on flights, negotiating with potential sellers of component services identified in the option set. Often, this will involve parallel negotiation in multiple marketplaces, and the use of different trading mechanisms (e.g. exchanges, auctions, RFQs). Furthermore, it may involve the negotiation of multiple complex parameters, for example: pricing policy, interaction processes, time constraints, and payment procedure.

One of the key problems Freightmixer faces is that of making commitments when negotiating simultaneously with customers and suppliers. It would like a scenario which avoids it making commitments to sell a service which it may not be able to deliver, or to buy a service it may not need. Hence it favours a scenario such as:

1. Freightmixer negotiates a price with a customer. The customer agrees to definitely buy the service, but Freightmixer doesn't commit to providing it.

2. Based on this known sale price, Freightmixer negotiates with several potential providers of component services. It agrees deals to maximise its profit, and commits to those deals.
3. Freightmixer returns to the customer, and commits to the original deal.

Notice that this scenario requires the customer to commit to an uncertain deal, to allow Freightmixer to avoid risk. A similar, dual, scenario where the component service providers commit to a deal without Freightmixer also committing will also give this. However, neither of these scenarios can be relied on. Firstly, buyers (resp. suppliers) may not want to use such a scenario as it places risk on them. Secondly, many market mechanisms (such as auctions) require commitment if Freightmixer is to negotiate in them.

Hence a more likely scenario, based on the example above, is:

1. Freightmixer observes the RFQ auction to ship a crate to SFO, and estimates (based on prior experience) what offer it would need to place to be likely to win, and the risk of losing associated with such an offer.
2. Based on the income it would receive, and taking into account the risk associated with the chance of losing, Freightmixer determines the maximum it is prepared to spend on procuring the composite service.
3. Using this as an upper bound, it places bids in the auctions for some of the component services. (eg space on the flights from LHR to BOS and BOS to SFO.) It must do this in such a way to procure an appropriate set of components for the cheapest price. In §4, we discuss this problem in more detail, and present a solution.
4. If it wins these auctions, it returns to the RFQ auction and places an offer there.

The outcome of this second phase is a series of contracts with suppliers that make sure that Freightmizer can deliver a composite service to the customer, and a contract with the end customer.

3.3 Service Execution

When FreightMixer has bought appropriate components to meet a customers need, and successfully negotiated with the customer to agree a contract, service can be delivered. During the service delivery phase, FreightMixer must ensure that the hand-over of the good at each stage of the journey takes place smoothly and appropriate paperwork is carried out.

A detailed discussion of the problems deriving from this phase is beyond the scope of this paper. For more information see (Salle et al., 2001; Morciniec et al., 2001).

In this way, FreightMixer is able to provide the same functionality as a large company despite the fact that its only assets are its market knowledge and organization ability. Using this expertise, it can compete with established transport companies with large infrastructure. We believe that service composition will play an important role in the future of e-commerce.

4 A Negotiation Algorithm Specification to Purchase Service Bundles

We now turn our attention to one specific aspect of the service composition problem - that of negotiating to purchase composable services. For the purposes of this initial work, we make certain restrictions on the scenario discussed above. Firstly, we assume that the service composer is buying from a set of auctions only. Secondly, we assume the customer of the service composer is offering to pay a certain fixed price. Hence, initially, we ignore the issue of simultaneously negotiating with the customer. We hope to address this in future work.

We consider how an agent involved in service composition should behave when participating in a set of auctions. It's aim is to buy a set of services which can be composed to sell on as a bespoke composite service, possibly to a specific customer with special requirements. There may be several ways of creating this composite service out of individual services for sale in the auctions. The agent's task is to purchase one such set of services which can be composed, without accidentally purchasing additional, unnecessary services.

In this section, we propose a possible approach for doing this. We first present the decision problem the agent is faced with, and then present the specification of an algorithm to perform service composition in this environment.

4.1 Specification of the Decision Problem

We assume our agent is participating in a set of auctions, \mathcal{A} . These auctions all start at roughly the same time, but may finish at different times. Each auction is selling one good or service. The auctions are English auctions with a fixed closing time. Participants can place bids at any time, provided the new bid is a minimum increment, ϵ , above the last bid. We choose units in which this minimum increment is 1. At the closing time, the good or service for sale is sold to the highest bidder at the price they bid. We also assume that the bid increment of each auction is very small with respect to the value of the good or service for sale.

To each subset $A \subset \mathcal{A}$ we associate a number $v(A)$, the "value" to the agent of winning the auctions A . By structuring the valuation of the agent as a function $v : 2^{\mathcal{A}} \rightarrow \mathbb{R}$, we allow for complements and substitutes in the normal fashion. We define a *bid set* to be a pair (A, \mathbf{p}) , where $A \subset \mathcal{A}$ and $\mathbf{p} : A \rightarrow \mathbb{Z}$ is a price function. The "utility" to the agent of winning a bid set (A, \mathbf{p}) is

$$u(A, \mathbf{p}) = v(A) - \sum_{a \in A} \mathbf{p}(a).$$

Our agent maintains a probabilistic model of the expected outcomes of each auction, based on past performance of similar auctions. (Discussion of some possible ways of generating this model is provided in (Preist et al., 2001a).)

To each auction $a \in \mathcal{A}$ is associated a price distribution $P_a : \mathbb{Z} \rightarrow [0, 1]$ representing the belief that, with probability $P_a(p)$, auction a will close at price p . We set $F_a(p) = \sum_{p' \geq p} P_a(p')$: the agent's believed probability that auction a will close at or above price p . For subsets $A \subset \mathcal{A}$ we define $P_A(\mathbf{p})$ to be the believed probability that the auctions in A will close at the prices specified by a *price function* $\mathbf{p} : A \rightarrow \mathbb{Z}$:

$$P_A(\mathbf{p}) = \prod_{a \in A} P_a(\mathbf{p}(a)), \quad (1)$$

and likewise F_A , the probability that that the auctions in A will close at or above the prices specified by \mathbf{p} :

$$F_A(\mathbf{p}) = \prod_{a \in A} F_a(\mathbf{p}(a)). \quad (2)$$

If the price in auction a is q , then the agent believes that the probability of a bid at price $p \geq q$ winning is

$$P_{win}(a, p, q) := \frac{P_a(p)}{F_a(q)}. \quad (3)$$

Similarly, for a collection of auctions A with current prices $\mathbf{q} : A \rightarrow \mathbb{R}$, the probability of the auctions closing at prices \mathbf{p} is

$$P_{win}(A, \mathbf{p}, \mathbf{q}) = \frac{P_A(\mathbf{p})}{F_A(\mathbf{q})}. \quad (4)$$

4.2 Specification of the algorithm

We now consider how the agent can use these beliefs to calculate information about expected future utility of deals it may win. Firstly, we define the notion of the expected utility $E(B, A, \mathbf{q})$ of a set of auctions B , given a set of observed prices \mathbf{q} , and given that the agent holds active bids in auctions A .

$$E(B, A, \mathbf{q}) = v(B) - C(B \cap A, \mathbf{q}) - C(B \setminus A, \mathbf{q} + 1) \quad (5)$$

where the function $C(S, \mathbf{q}')$ is the expected cost of winning the auctions S at prices greater than or equal to \mathbf{q}' :

$$C(S, \mathbf{q}') = \sum_{\mathbf{p}' \geq \mathbf{q}'} \sum_{a \in S} P_{win}(a, \mathbf{p}'(a), \mathbf{q}'(a)) \mathbf{p}'(a) \quad (6)$$

The expected utility of a set of auctions is thus the value of the bundle, minus the expected cost of winning each of the auctions. The latter is calculated by using the believed probability that the auction will finish at each given price, if our agent places a bid at that price. We restrict $p' > q$ for auctions $B \setminus A$, in (5) because we know that the agent does not hold bids in these auctions at prices \mathbf{q} , and so has no probability of winning at these prices.

The expression (5) gives us some idea of the intrinsic value of a bundle of goods B , but is not the expected return for placing a single bid in the auctions in B . In general such a bid does not *have* an expected return: we must reason over complete strategies.

Consider the expected value (given that prices are currently \mathbf{q} , and the agent holds the active bids A) of the following strategy, which we call “commitment to B ”: The agent chooses a set of auctions B , and for all future time steps, will always bid on any elements of B in which it does not hold active bids. If the agent sticks to this commitment, then we know its future choices, and so precise formulae for expected return can be calculated.

Let S be a possible set of auctions that the agent may win using this strategy: $B \subset S \subset A \cup B$. The probability that the auctions $S \setminus B$ will not be outbid, while the auctions $A \setminus S$ are, is

$$P_{ret}(S, A, \mathbf{q}) = \frac{F_{A \setminus S}(\mathbf{q} + 1) P_{S \setminus B}(\mathbf{q})}{F_{A \setminus B}(\mathbf{q})}$$

Given this eventuality, the expected utility is evaluated in the same way as (5), except that instead of $v(B)$, the value we obtain is $v(S)$, and we incur additional costs for each auction in $S \setminus B$ that we win.

It follows that the expected value for following the commitment to B is

$$E_c(B, A, \mathbf{q}) = E(B, A, \mathbf{q}) + \sum_{B \subset S \subset A \cup B} P_{ret}(S, A, \mathbf{q}) \left((v(S) - v(B)) - \sum_{a \in S \setminus B} \mathbf{q}(a) \right) \quad (7)$$

The terms in this expression for which $S = B$ are the desired outcomes. The other terms correspond to obtaining some non-empty collection $S \setminus B$ of goods that do not contribute to our desired bundle B . Although they could still provide positive value, it is anticipated that in the service composition arena, where goods tend to complement one another, the slight increase of $v(S)$ with respect to $v(B)$ will not be large enough to compensate for the increase in costs $\sum_{a \in S \setminus B} \mathbf{q}(a)$, and each of these terms would have a negative impact on the expected value of the commitment.

The algorithm (COMPOSER) we propose is that at each time step the agent calculates the commitment B which has largest expected utility $E_c(B, A, \mathbf{q})$ given the currently held bids A and prices \mathbf{q} , and places the minimal bids required to take the lead in $B \setminus A$. In practice, this means it will bid initially in the auctions which have highest a-priori expected utility. It will continue to compete in these auctions, placing more bids when outbid. However, if sufficient competing bids are placed to reduce the expected utility of this set of auctions, then it may change to another set of auctions, for another bundle. It will do this if the expected gain from changing to this new bundle outweighs the expected cost of currently held bids which appear in the old bundle but not in the new bundle.

There are two obvious problems with this algorithm:

- By its very nature, our algorithm does not in fact commit, since it re-evaluates its options at each opportunity. However, the value $E_c(B, A, \mathbf{q})$, which is truly the expected value of committing to bid on B , and hence is *not* the expected value according to the specified algorithm, is none-the-less (we claim) a good indication of the optimal choice to make. The estimate we use is conservative, in that the agent chooses a single bundle that will give the best overall expected utility. Choosing a different bundle for each possible outcome can only improve on this. We have adopted this approach initially, as we believe that it will provide good performance in the majority of situations. Experimentation and further analysis will be necessary to test this hypothesis.
- In practice, if the number of auctions is large, it will be difficult to evaluate (7) given realistic computational resource bounds. Ideally, if we had perfect information and unlimited computation time, we would calculate this accurately. Finding appropriate simplifications which still give good results is a topic to which we will return in further work.

4.3 Worked Example

To illustrate how this analysis operates, we return to the FreightMixer scenario described in §3. Based on past histories of similar auctions to the ones that were selected during the matchmaking phase, Freightmixer creates beliefs about the expected distribution of closing prices of these auctions. We assume that the closing prices are uniformly distributed over the following sets

$$\begin{aligned}
a & : \{40, 45, \dots, 135, 140\} \\
b & : \{20, 25, \dots, 95, 100\} \\
c & : \{130, 135, 140, 145, 150\} \\
d & : \{50, 55, \dots, 105, 110\} \\
e & : \{80, 85, \dots, 115, 120\} \\
f & : \{30, 35, \dots, 65, 70\}
\end{aligned} \tag{8}$$

Before bidding begins, the agent holds no bids. We assume that the current price function \mathbf{q}_0 lies just below all of the above prices. The expected utilities of committing to each of the bundles which we seriously consider are therefore the same as the expected values of the bundles:

$$\begin{aligned}
E(\{a, b\}, \emptyset, \mathbf{q}_0) & = 50 \\
E(c, \emptyset, \mathbf{q}_0) & = 60 \\
E(\{d, f\}, \emptyset, \mathbf{q}_0) & = 70 \\
E(\{e, f\}, \emptyset, \mathbf{q}_0) & = 50
\end{aligned} \tag{9}$$

The agent therefore chooses to bid in $\{d, f\}$, even though the bundle $\{a, b\}$ has greater initial value¹. This can be seen as sensible, given that by bidding for $\{a, b\}$ the agent runs the risk of ending up committed to this bundle, even though it is non-optimal in expectation. It can be argued that the risk of such a commitment is low, since if the prices in a or b become prohibitively high, then the agent can simply wait, and with high probability will be outbid in these auctions, and so de-committed from them. It follows, however, that the agent also has correspondingly low probability of winning $\{a, b\}$ at these low prices: It is precisely the payoff between the chance of a good deal and the chance of being committed to a bundle which our algorithm seeks to address.

This payoff between expected return and commitment explicitly comes into effect in this example if prices in d or f rise without the auctions closing.

Suppose that prices in c , d and f have risen to 140, 75, 40. If the agent holds no bids, it should bid in c , since the expected cost of c given these prices is 147.5, whereas the expected cost of $\{d, f\}$ is $97.5 + 57.5 = 155$. Even if the agent holds the active bid in auction f , then the expected loss from accidentally winning f , which is the cost of f , 40, multiplied by the probability that no other agent will bid in f , 0.16666, is less than the difference between the expected costs of c and $\{d, f\}$, and so it is still preferable to bid for c despite the risk of winning f .

If, on the other hand, the agent holds the leading bid in d , then the potential cost of winning d is too high (expected loss 13.333) to risk bidding for c , despite the fact that c is, by now, expected to do better than $\{d, f\}$.

The benefit of this algorithm over a greedy one is clear in the same situation: A greedy algorithm would continue to pursue $\{d, f\}$ in preference to c until its current aggregate price was as large as that in c . For example, if prices in c , d and f were 145, 95 and 45 (and the agent held no bids) then a greedy algorithm would bid in $\{d, f\}$ in preference to c . The reason why this is foolish is that the probability of winning one of these auctions but not the other (at these prices), is large: 40%. If the agent wins one, then it is committed to bidding for the other, despite its large expected cost.

¹The bid-price for the bundle $\{a, b\}$ is 60, much lower than that of $\{d, f\}$, at 80.

5 Related Work

Research into automated negotiation has long been an important part of distributed AI and multi-agent systems. Initially it focused primarily on negotiation in collaborative problem solving, as a means towards improving coordination of multiple agents working together on a common task. (Laasri et al., 1992) provide an overview of the pioneering work in this area. As electronic commerce became increasingly important, the work expanded to encompass situations with agents representing individuals or businesses with potentially conflicting interests. The Contract Net (Smith, 1980) provides an early architecture for the distribution of contracts and subcontracts to suppliers. It uses a form of distributed request-for-proposals. However, it does not discuss algorithms for determining what price to ask in a proposal. (Jennings et al., 1996) use a more sophisticated negotiation protocol to allow the subcontracting of aspects of a business process to third parties. This is primarily treated as a one-to-one negotiation problem, and various heuristic algorithms for negotiation in this context are discussed in (Faratin et al., 1998). (Vulkan and Jennings, 1998) recast the problem as a one-to-many negotiation, and provide an appropriate negotiation protocol to handle this.

Other relevant work in one-to-one negotiation includes the game-theoretic approach of (Rosenschein and Zlotkin, 1994) and the logic-based argumentation approach of (Parsons et al., 1998). As much electronic commerce involves one-to-many or many-to-many negotiation, the work in the agent community has broadened to explore these cases too. The Michigan AuctionBot (Wurman et al., 1998) provides an automated auction house for experimentation with bidding algorithms. The Spanish Fishmarket developed by (Rodríguez-Aguilar et al., 1997) provides a sophisticated platform and problem specifications for comparison of different bidding strategies in a Dutch auction, where a variety of lots are offered sequentially. The Kasbah system (Chavez et al., 1997) featured agents involved in many-to-many negotiations to make purchases on behalf of their users. However, the algorithm used by the agents (a simple version of those in (Faratin et al., 1998)) was more appropriate in one-to-one negotiation, and so gave rise to some counter-intuitive behaviours by the agents. (Preist and van Tol, 1998) and (Cliff and Bruten, 1998) present adaptive agents able to effectively bid in many-to-many marketplaces, and are the first examples of work which borrow techniques from experimental economics to analyze the dynamics of agent-based systems. (Preist, 1999) demonstrates how these can be used to produce a market mechanism with desirable properties. (Park et al., 1998) and (Park et al., 1999) present a stochastic-based algorithm for use in the University of Michigan Digital Library, another many-to-many market.

(Gjerstad and Dickhaut, 1998) use a belief-based modeling approach to generating appropriate bids in a double auction. Their work is close in spirit to ours, in that it combines belief-based learning of individual agents bidding strategies with utility analysis. However, it is applied to a single double auction marketplace, and does not allow agents to bid in a variety of auctions. (Vulkan and Preist, 1999) use a more sophisticated learning mechanism that combines belief-based learning with reinforcement learning. Again, the context for this is a single double auction marketplace. Unlike Gjerstad's approach, this focuses on learning the distribution of the equilibrium price. The work of (Garcia et al., 1998) is clearly relevant. They consider the development of bidding strategies in the context of the Spanish Fishmarket tournament. Agents compete in a sequence of Dutch auctions, and use a combination of utility modeling and fuzzy heuristics to generate their bidding strategy. Their work focuses on Dutch rather than English auctions, and on a sequence of auctions run by a single auction house rather than parallel auctions run by

multiple auction houses.

In our previous work, (Preist et al., 2001a), we have presented algorithms which allow agents to participate simultaneously in multiple auctions for the purchase of a number of similar goods. In (Preist et al., 2001b), we show how agents using these algorithms in multiple auctions can create a more efficient and stable market. It is interesting to contrast our analysis with that of (Greenwald and Kephart, 1999). They demonstrate that the use of dynamic price-setting agents by sellers, to adjust their price in response to other sellers, can lead to an unstable market with cyclical price wars occurring. We, however, show that (in a very different context) the use of agents improves the dynamics and stability of the market. From this, we can conclude that agent technology is not a-priori ‘good’ or ‘bad’ for market dynamics, but that each potential role must be studied to determine its appropriateness.

In this paper, we have extended our earlier work to develop algorithms to purchase heterogeneous bundles of goods from multiple auctions. An alternative approach is to attempt to provide the right market mechanism in the first place, providing a centralized point of contact for all buyers and sellers to trade. (Sandholm, 2000) proposes a sophisticated marketplace able to handle combinatorial bidding, and able to provide guidance to buyers and sellers as to which market mechanism to adopt for a particular negotiation. In the long term, as the different auction houses merge or fold and only a few remain, this approach will be ideal. In the short term, we expect improved market dynamics will occur through autonomous agents in multiple auctions.

6 Conclusions and Future Work

In the future, service composition will play an essential role in e-commerce. Composite service products will be created on the fly in response to customer requests. However, if this is to happen, several technical problems must be overcome.

We have focused on the key problem of effective negotiation for service composition, and presented the specification of an algorithm to perform this task. The algorithm competes in multiple simultaneous auctions, placing appropriate bids to create service bundles. In future work we will focus on the two areas mentioned as problematic at the end of §4.2.

Firstly, we will pursue formal methods to calculate the expected value of a specified bidding algorithm. We already have a closed form (7) for the value of the scale of strategies described as “committed”, but not (for example) for *COMPOSER*. By finding expressions which bound, or better still equal (in expectation) the value that can be extracted from a given algorithm, we can be sure that a given algorithm has theoretical performance properties which otherwise we would have to guess at, or simulate.

On the other hand, theoretical models for algorithm effectiveness are useless if those algorithms are impractical, and so another main focus of our future work in this area will be finding algorithms which are of low complexity. There is likely to be a payoff between complexity and optimality; it is our goal to develop practical algorithms which have high, well understood value.

Throughout the paper we have made assumptions regarding risk which are often not appropriate in practice. In particular we have assumed risk-neutrality, whereas in reality many potential users of service aggregation algorithms are risk-averse. In future work we will extend the algorithmic analysis we have just described to cases involving alternative

risk profiles. We also plan to address the problem of simultaneous negotiation not only with the suppliers, but also with the purchasers. We would also like to generalize the work beyond the English auction, and to allow auctions to have staggered opening times.

References

- Chavez, A., Dreilinger, D., Guttman, R., and Maes, P. (1997). A real-life experiment in creating an agent marketplace. In *Proc. 2nd Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Systems*.
- Cliff, D. and Bruten, J. (1998). Less than human: Simple adaptive trading agents for CDA markets. In *Proc. of the 1998 Symposium on Computation in Economics, Finance, and Engineering: Economic Systems*.
- Dale, J. and Mamdani, E. (2001). Open standards for interoperating agent based systems. forthcoming.
- Dignum, F. and Weigand, H. (1995). Modelling communication between cooperative systems. In *Proc. 7th Conf. on Advanced Information Systems Engineering (CAISE'95)*, pages 140–153.
- Faratin, P., Sierra, C., and Jennings, N. (1998). Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, 3-4(24):159–182.
- Finin, T. and Fritzson, R. (1994). KQML as an agent communication language. In *Proc. Third International Conference on Information and Knowledge Management (CIKM'94)*. ACM Press.
- Garcia, P., Gimenez, E., Godo, L., and Rodriguez-Aguilar, J. (1998). Possibilistic-based design of bidding strategies in electronic auctions. In *Proc. 13th Biennial European Conference on Artificial Intelligence*.
- Gjerstad, S. and Dickhaut, J. (1998). Price formation in double auctions. *Games and Economic Behaviour*, 22(1):1–29.
- Greenwald, A. R. and Kephart, J. O. (1999). Shopbots and pricebots. In *Proc. 16th Int. Joint Conf. on Artificial Intelligence*.
- Jennings, N. R., Faratin, P., Johnson, M. J., O'Brien, P. O., and Wiegand, M. E. (1996). Using intelligent agents to manage business processes. In *First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM-96)*, pages 345–360.
- Laasri, B., Laasri, H., Lander, S., and Lesser, V. (1992). Generic model for intelligent negotiating agents. *International Journal of Intelligent and Cooperative Information Systems*, 2(1):291–317.
- Morciniec, M., Salle, M., and Moynihan, B. (2001). Towards regulating electronic communities with contracts. In *Proc. 2nd Workshop on Norms and Institutions in MAS, at Autonomous Agents 2001*.
- Norman, T. J. and Reed, C. (2000). Delegation and responsibility. In *Proc. 3rd UK Workshop on Multi-Agent Systems*.

- Park, S., Durfee, E., and Birmingham, W. (1998). Emergent properties of a market-based digital library with strategic agents. In *Proc. 2nd Int. Conf. on Multi-Agent Systems*.
- Park, S., Durfee, E., and Birmingham, W. (1999). An adaptive agent bidding strategy based on stochastic modelling. In *Proc. 3rd Int. Conf. on Autonomous Agents*.
- Parsons, S., Sierra, C., and Jennings, N. R. (1998). Agents that reason and negotiate by arguing. *Journal of Logic and Computation*, 8(3):261–292.
- Peters, T., Austin, K., and Peters, T. J. (1984). *A Passion for Excellence : The Leadership Difference*. Warner Books.
- Piccinelli, G. and Lynden, S. (2000). Concepts and tools for e-service development. In *Proc. 7th Int. Workshop HP OVUA*.
- Piccinelli, G. and Mokrushin, L. (2001). Dynamic e-service composition in DySCo. part of the 21st Int. Conf. on Distributed Computing Systems (ICDCS-21).
- Preist, C. (1999). Commodity trading using an agent-based iterated double auction. In *Proc. 3rd Int. Conf. on Autonomous Agents*.
- Preist, C., Bartolini, C., and Philips, I. (2001a). Algorithm design for agents which participate in multiple simultaneous auctions. In Dignum, F. and Cortes, U., editors, *Agent Mediated Electronic Commerce III*, Lecture Notes in AI. Springer Verlag.
- Preist, C., Byde, A., and Bartolini, C. (2001b). Economic dynamics of agents in multiple auctions. In *Proc. 5th Int. Conf. on Autonomous Agents*.
- Preist, C. and van Tol, M. (1998). Adaptive agents in a persistent shout double auction. In *Proc. 1st Int. Conf. on the Internet, Computing and Economics*. ACM Press.
- Rodriguez-Aguilar, J., Noriega, P., Sierra, C., and Padget, J. (1997). A Java-based electronic auction house. In *Proc. Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Systems*, pages 207–224.
- Rosenschein, J. and Zlotkin, G. (1994). *Rules of Encounter*. MIT Press.
- Salle, M., Morciniec, M., and Moynihan, B. (2001). A conceptual architecture for market governance in trusted electronic marketplaces. In *Proc. 2nd Workshop on Norms and Institutions in MAS, at Autonomous Agents 2001*.
- Sandholm, T. (2000). eMediator: A next generation electronic commerce server. In *Proc. 4th Int. Conf. on Autonomous Agents*.
- Smith, R. G. (1980). The Contract Net protocol: High level communication and control in a distributed problem solver. *IEEE Transactions on Computing*, 12(29):1104–1113.
- Tan, Y. and Thoen, W. (1999). A logical model of directed obligations and permissions to support electronic contracting. *International Journal of Electronic Commerce*, 3(2):87–104.
- Vulkan, N. and Jennings, N. (1998). Efficient mechanisms for the supply of services in multi-agent environments. In *Proc. 1st Int. Conf. on the Internet, Computing and Economics*. ACM Press.

- Vulkan, N. and Preist, C. (1999). Automated trading in agents-based markets for communication bandwidth. In *Proc. 2nd UK Workshop on Multi-Agent Systems*.
- Wurman, P., Wellman, M., and Walsh, W. (1998). The Michigan internet AuctionBot: A configurable auction server for human and software agents. In *Proc. 2nd Int. Conf. on Autonomous Agents*.