

On Agent-Mediated Electronic Commerce

Minghua He, Nicholas R. Jennings, and Ho-Fung Leung

Abstract—This paper surveys and analyzes the state of the art of agent-mediated electronic commerce (e-commerce), concentrating particularly on the business-to-consumer (B2C) and business-to-business (B2B) aspects. From the *consumer buying behavior* perspective, agents are being used in the following activities: need identification, product brokering, buyer coalition formation, merchant brokering, and negotiation. The roles of agents in B2B e-commerce are discussed through the *business-to-business transaction* model that identifies agents as being employed in partnership formation, brokering, and negotiation. Having identified the roles for agents in B2C and B2B e-commerce, some of the key underpinning technologies of this vision are highlighted. Finally, we conclude by discussing the future directions and potential impediments to the wide-scale adoption of agent-mediated e-commerce.

Index Terms—Agent-mediated electronic commerce, intelligent agents.

1 INTRODUCTION

ELECTRONIC commerce (e-commerce) is increasingly assuming a pivotal role in many organizations. It offers opportunities to significantly improve (make faster, cheaper, more personalized, and/or more agile) the way that businesses interact with both their customers and their suppliers. However, in order to harness the full potential of this new mode of commerce, a broad range of social, legal, and technical issues need to be addressed. These issues relate to things such as security, trust, payment mechanisms, advertising, logistics, and back office management [152], [45], [138], [169].

Even more fundamental than these issues, however, is the very nature of the various actors that are involved in e-commerce transactions. In most current (first generation) e-commerce applications, the buyers are generally humans who typically browse through a catalog of well-defined commodities (e.g., flights, books, compact discs, computer components) and make (fixed price) purchases (often by means of a credit card transaction). However, this modus operandi is only scratching the surface of what is possible. By increasing the degree and the sophistication of the automation, on both the buyer's and the seller's side, commerce becomes much more dynamic, personalized, and context sensitive. These changes can be of benefit to both the buyers and the sellers. From the buyer's perspective, it is desirable to have software that could crawl all the available outlets to find the most suitable one for purchasing the chosen good (e.g., the one that offers the cheapest price, the highest quality, or the fastest delivery time) and that could then go through the process of actually purchasing the good, paying for it, and arranging delivery at an appropriate time. From a seller's perspective, it is desirable to have software that could vary its offering (in

terms of price, quality, warranty, and so on) depending on: the customer it is dealing with (e.g., offering discounts or special offers to particular target groups), what its competitors are doing (e.g., continuously monitoring their prices and making sure its own price is competitive), and the current state of its business (e.g., if it has plenty of a particular item in stock, it may be appropriate to reduce the price in order to try and increase demand).

To achieve this degree of automation and move to second generation e-commerce¹ applications, we believe that a new model of software is needed. This model is based upon the notion of *interacting agents* [64] (hence, the term "agent-mediated electronic commerce"). Against this background, the aim of this paper is to motivate the use of agents in e-commerce, to highlight the roles that agents can and are fulfilling in this domain, and to review some of the key technologies that underpin this vision. First, however, we define exactly what we mean by the terms "agent" and "electronic commerce."

1.1 Interacting Agents

An agent is a software program that acts flexibly on behalf of its owner to achieve particular objectives [64].² To do this, the software must exhibit the following properties [162]:

- it needs to be *autonomous*: capable of making decisions about what actions to take without constantly referring back to its user;
- it needs to be *reactive*: able to respond appropriately to the prevailing circumstances in dynamic and unpredictable environments;

1. Second generation systems are characterized here as having a greater degree of automation on both the buyer's and the seller's side. Like many classifications, however, this distinction is not absolute and there are areas of uncertainty between the generations. Moreover, the same is also true of agent-mediated e-commerce in general. While there are some systems that are clearly agent-mediated and some that are not all agent-based, there is a degree of uncertainty in some cases. This is caused by the fact that in such systems, agents are rarely the only technology that is used. Often, an e-commerce system will be composed of a variety of technologies, only a fraction of which will be agent-based.

2. Agents are not just used in the domain of e-commerce, although this is arguably their most popular domain. Rather, agent technology should be viewed as a general solution paradigm for developing complex systems [63]. Overviews of the application of agents in other domains can be found in [23], [66], [69], [109].

- M. He and N.R. Jennings are with the Department of Electronics and Computer Science, University of Southampton, Southampton SO17 1BJ, United Kingdom. E-mail: {mh00r, nrj}@ecs.soton.ac.uk.
- H.-F. Leung is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong SAR, P.R. China. E-mail: lhf@cse.cuhk.edu.hk.

Manuscript received 7 Feb. 2002; revised 28 Aug. 2002; accepted 25 Nov. 2002.

For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number 115866.

- it needs to be *proactive*: able to act in anticipation of future goals so that its owner's objectives are met.

Thus, for example, a buyer may instruct their agent to: 1) find a reasonably cheap notebook computer, with the latest technical specification, that can be delivered within a week and that has a two year warranty, or 2) to find a flight that takes her from London to San Francisco with a weekend stop over in New York. Similarly, a seller may instruct an agent to: 1) monitor the prices of all its known competitors and automatically adjust its offerings (either up or down) so that they remain attractive to their target audience, or 2) to offer a reasonable discount scheme and better credit facilities to highly valued customers.

While pursuing their objectives, the agents will invariably need to *interact* with other similarly autonomous agents. This interaction can vary from simple communication (e.g., a buyer agent asks a seller agent how much a particular computer costs), to more elaborate forms of social interaction (e.g., cooperation, coordination, and negotiation). In the latter case, for example, an agent may be required to:

- participate in an online auction, e.g., monitoring bids, making bids, and withdrawing from the auction;
- negotiate on behalf of its owner, e.g., to ensure that the desired good will be delivered in time or to make the price acceptable;
- cooperate with other agents, e.g., two sellers of the same product may need to pool their resources in order to meet a large customer order, or two sellers may bundle their distinct (complementary) offerings to make a single more desirable product.

In addition to the buyers and the sellers, the third key actor in agent-mediated e-commerce applications is the *market owner*. The owner is the individual or organization that controls the environment (sets the rules) in which buyers and sellers trade. In first generation systems, the market owner is usually synonymous with the seller. However, this need not always be the case; examples of other possibilities are the numerous third party auction sites that now exist (e.g., eBay (<http://www.ebay.com>), Fastparts (<http://www.fastparts.com>) and Freemarkets (<http://www.freemarkets.com>) and situations where buyers put out requests for tender (e.g., Labx (<http://www.labx.com>) and General Electric (<http://www.ge.com>)). Generally speaking, there will be multiple e-markets trading in different types of goods (e.g., e-markets for purchasing holidays, e-markets for buying computer equipment, and e-markets for finding plumbers). Moreover, there will, in many cases, be multiple e-markets for the same (or similar) goods (i.e., there will not be a single market for dealing with holidays). Both within and between different vertical market segments, there will be a significant variety in the way that e-markets are organized. These will vary from simple, fixed-price catalogs, through various forms of online auction, to sites where buyers and sellers can negotiate directly with one another.

1.2 Electronic Commerce

A general definition of e-commerce, given by the Electronic Commerce Association, is: "electronic commerce covers any

form of business or administrative transaction or information exchange that is executed using any information and communications technology" [147]. However, we believe this definition is too broad and so we limit this paper to covering commercial activities conducted on the Internet [56]. Therefore, other forms of remote transactions (e.g., ordering an air ticket over the telephone or buying a computer by credit card) are not considered here.

According to the nature of the transactions, the following types of e-commerce are distinguished [152]: business-to-business (B2B), business-to-consumer (B2C), consumer-to-consumer (C2C), consumer-to-business (C2B), nonbusiness e-commerce (use of the Internet by nonbusiness organizations such as academic institutions or government agencies to reduce expenses or improve services), and intrabusiness e-commerce. Currently, however, most applications are either B2C or B2B and, therefore, these are the two areas that we focus on here. In more detail, B2C mainly refers to online retailing transactions with individual customers, where shoppers can conduct transactions through a company's homepage. B2B refers to the transactions where both sellers and buyers are business corporations. Although most of the initial Web-based e-commerce was in the B2C domain, B2C now constitutes a smaller portion of the overall landscape. For example, B2B transactions are expected to be in the range of \$800 billion by 2003, which is five times as much as B2C [132]. Moreover, it is widely believed that B2B will be the predominant means of doing business within the next five years [141], [133].

Within this context, e-commerce systems provide both commercial information (such as products' prices and available quantities) and facilitate various commercial actions (e.g., buying, selling, and negotiation). The increasing use of information technology in this area has led to fundamental changes in the way these commercial activities are undertaken (e.g., the rise of dynamic pricing, the ability to easily compare many goods, and the ability to negotiate contracts much more frequently) [133] and as more advanced information processing techniques are brought to bear, so the potential for further changes increases.

1.3 Agents for Electronic Commerce

IDC (<http://www.idc.com>) estimates that the global market for software agents grew from \$7.2 million in 1997 to \$51.5 million in 1999, and that it will reach \$873.2 million in 2004, with a compound annual growth rate of 76.2 percent between 1999 and 2004. They also assume that the dramatic growth in B2B e-commerce will accelerate the demand for agents. To this end, and in order to motivate the potential of agent-mediated e-commerce, we consider the following medium term scenarios³ as examples of what will be possible [68].

Scenario 1: Finding closest match to buyer's requirements. A buyer decides that they would like a holiday in one of the Greek islands, they would like to go next Friday, they would

3. We do not focus on current applications because they do not adequately highlight the full potential of agent-mediated e-commerce. Current applications tend to use agents in reasonably straightforward ways. Also, organizations that have adopted agent-based techniques often do not disclose this fact for reasons of retaining competitive advantage. Focusing on medium-term scenarios overcomes both of these concerns without having to gaze too far into the future (which is notoriously unpredictable).

like to fly from London, and that the total cost should be less than 300 pounds. Their software agent is instructed to go and find out what is available and to report the options back to the user who will make the ultimate choice. In order to fulfill this objective, the buyer agent determines those e-markets that deal with leisure activities. From those, it tries to find out holidays that meet the specified requirements. However, it finds no appropriate fixed price offerings and after observing the outcome of several online auctions, it decides that it will be very unlikely that it will be able to meet all of these requirements. It therefore decides to relax some of the user's constraints and tries to find holidays that are similar. The agent decides to relax the user's stated requirements in the following way: it looks for holidays to the Greek islands that leave any day next week, that leave from non-London airports in the United Kingdom next Friday, and that cost up to 400 pounds. With these new requirements in place, the buyer agent returns to the relevant e-markets, collects the offerings that satisfy these new requirements, and returns them to its user with an explanation of why it acted in this way.

Scenario 2: Acting across multiple e-markets. A buyer decides that they would like to purchase a new laptop computer; they want a reasonably high specification, are prepared to pay for a good quality brand name, but it must be delivered within a week. Their software agent is instructed that they are prepared for the agent to find the most appropriate model, negotiate the best potential deal available, but that the user would like to make the final choice about purchase. In order to fulfill this objective, the buyer agent determines those e-markets that deal with selling computer equipment. From these, it selects those e-markets that offer products that meet the user's specification. In order to determine those machines that fit the specification, the buyer agent examines the sites of a number of computer manufacturers to determine the latest specification information and to determine an approximate price to pay. Armed with this information, the agent formulates a strategy for making a deal. The agent knows the maximum price it needs to pay (this will be the minimum of the cheapest fixed price offerings that are available in the catalogs). From this baseline, the agent tries to negotiate directly with several of the suppliers to see if they are willing to reduce the price (or bring forward the delivery time). In parallel to this, the agent tracks a number of online auctions to see if the same good can be purchased more cheaply (it will not actually bid in the auctions since submitting a bid would constitute a commitment on behalf of its buyer). When it has completed its negotiations (or before if a very good deal appears in an auction), the buyer agent reports back a ranked list of purchasing options to its owner. The owner then makes their choice and instructs their agent to complete the deal (including arranging payment and setting the delivery time and place).

Scenario 3: Coalition formation. A bakery agent receives a request for tender from a supermarket agent who wishes to purchase 500 iced buns a day throughout the summer period. The bakery agent has sufficient capacity to make 300 buns per day. However, the bakery would like to set up links with the supermarket and so is keen to see if it can fulfill the order. Thus, rather than simply turning the order away, the bakery instructs its agent to search for a partner who will produce the remaining 200 buns for the rest of the

summer period. In order to achieve this, the bakery agent contacts all the other sellers present in e-markets that offer iced buns. The bakery agent indicates it has a demand for 200 buns per day for the summer period and asks whether any of the other bakeries would like to join in a partnership with it to meet the supermarket's need. A number of potential collaborators come forward. The bakery agent then conducts a series of negotiations with these agents in order to set up the terms and conditions of the partnership. Eventually, a deal is reached and the bakery agent reports details of the arrangements back to the bakery.

In terms of the nomenclature outlined in Section 1.1, scenarios 1 and 2 fall into the B2C domain. The former shows that agents can, on behalf of their owners, locate and retrieve information and make reasonable decisions (relaxing the constraints of the search) based on the owner's profile. The latter scenario demonstrates how agents negotiate with multiple suppliers, monitor multiple auctions, and use intelligent strategies to find the best deal for the users. The agents in the third scenario represent companies/organizations in a B2B context. This example not only shows how agents can collaborate with one another to achieve a common goal, but also shows how an agent selects the best partners through negotiation.

Although agents can be used in a closed loop fashion (i.e., without human intervention), in many cases users will simply not be willing to delegate complete autonomy to them. Moreover, the degree of automation that user's find acceptable is likely to vary between individuals and between tasks for the same individual. For example, some users will not want any automated support—they will directly enact all phases of the trade themselves. Others may be willing to use agents to collate information and present them with options from which they make the subsequent purchasing decision. While yet others, will be happy to delegate all trading activities to their software agents. To reflect this situation, Fig. 1 shows the range of the automation that a software agent may be given.

1.4 Structure of the Paper

This paper builds upon a number of previous reviews of agent-mediated e-commerce. Particularly prominent among these are Guttman et al.'s review of agents in B2C e-commerce [54] and Sierra and Dignum's roadmap of agent-mediated e-commerce in Europe [135]. Other important sources include [88], [166], [121], [8], [17], [71], [86], [90], [102]. Our paper both extends and updates this previous material and also tries to present a more integrated and coherent view on the field. For example, in contrast to the aforementioned work, we categorize and systematically analyze applications of agent-based e-commerce in the B2C and B2B domains (using the consumer buying behavior (CBB) model and the business-to-business transaction (BBT) model, respectively). We extend the traditional CBB model so that it covers more B2C behaviors (such as buyer coalition formation), and we identify more uses for agents in the BBT model.

In more detail, the rest of this paper is structured as follows: Section 2 describes the basic roles and techniques of agents in B2C e-commerce. Section 3 performs a similar analysis for B2B e-commerce. Section 4 explores some of the key underpinning technologies that are used to support B2C and B2B e-commerce. Finally, Section 5 concludes this

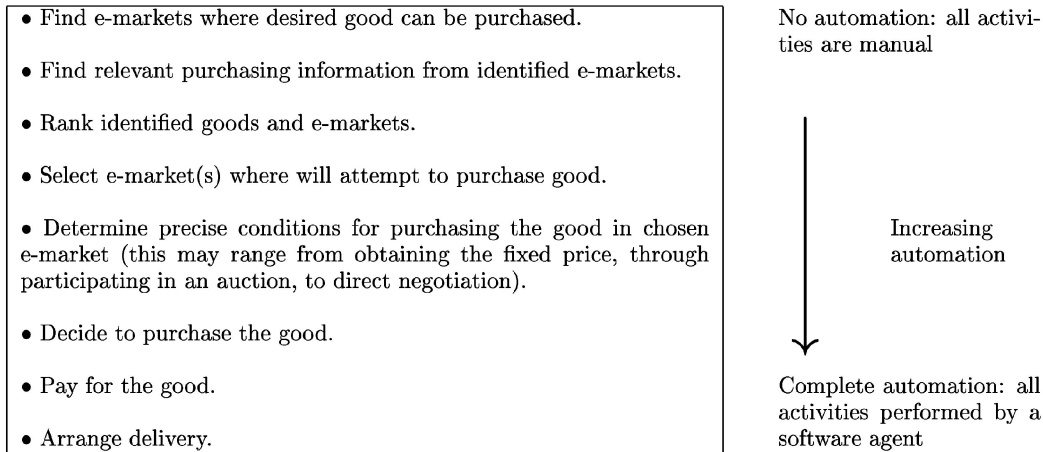


Fig. 1. Varying degree of agent automation.

paper and highlights some of the key challenges facing agent-mediated e-commerce.

2 AGENTS IN B2C E-COMMERCE

B2C e-commerce is becoming more widespread as more people come to recognize its convenience and its ability to offer a quick response to requests and as more products/services become available [102]. As this adoption spreads, the impetus for employing software agents increases in order to enhance and improve the trading experience. In order to systematically analyze the tasks that agents can assist with, we employ the CBB model (based on [54]) to capture consumer behavior (see Fig. 2). From the CBB model perspective, we believe agents can act as mediators in five of the stages: need identification, product brokering, buyer coalition formation, merchant brokering, and negotiation.⁴ Sometimes, the boundary between merchant brokering and negotiation is not always clear cut (because negotiation is sometimes also involved in brokering). For example, Jung and Jo [70] introduce a brokering technique that uses a negotiation protocol to match seller and buyer agents; in the brokering service of [12], a multiattribute auction is proposed to find a suitable supplier for a buyer; and in [32], the brokering service involves finding the optimal winner through a combinatorial auction. Against this background, each of the five above-mentioned agent-mediated stages is explored in more detail in the remainder of this section.

2.1 Need Identification

In this stage, the customer recognizes a need for some product or service. This need can be stimulated in many different ways (e.g., by advertisement, through friends, and

4. The sixth stage (purchase and delivery) involves paying for the transaction and arranging delivery of the goods/services. Here, the key problems are to ensure safe payment and delivery, problems that are common to e-commerce in general. The last stage involves product services (e.g., repair and upgrade services) and evaluation (measuring the degree of satisfaction of the user about the goods and the buying procedure). Generally speaking, however, these two stages have little that is specific to agent-mediated e-commerce and, thus, they are not discussed in detail here (interested readers can refer to [45], [152], [96], [140] for more information on these topics).

so on). However, in the agent-mediated e-commerce world, it can also be stimulated by the user's agent. Such an agent is typically called a notification agent.⁵ To do this, the notification agent needs to have a profile for the user. This profile can be obtained in many different ways: through observing the user's behavior [15], through direct elicitation techniques [114], or through inductive logic programming techniques [28]. Once the profile is installed in the agent, it can notify the user whenever an appropriate good/service becomes available (i.e., the user's profile matches a good/service catalog). For example, in Amazon Delivers (<http://www.amazon.com>), the latest reviews of exceptional new titles in categories that interest the user are sent automatically and Fastparts uses "AutoWatch" to allow users to list parts they need and notify them if those parts become available for sale.

2.2 Product Brokering

Having ascertained a need, the product brokering stage involves an agent determining what product to buy to satisfy this need. The main techniques used by the brokers in this stage are: feature-based filtering, collaborative filtering, and constraint-based filtering. Table 1 shows a number of exemplar e-commerce systems that exploit these techniques. Feature-based filtering involves selecting products based on feature keywords. For example, suppose a customer wants to buy a Sony notebook computer through Amazon. His agent selects the "Computers" category first, then indicates "Sony" in the *brands* field, and the notebook computers with these features are returned. Collaborative filtering [131] involves giving an agent personalized recommendations based on the similarities between different users' preference profiles. Here, the product rating of

5. This is a reasonably simple type of agent. It acts autonomously to inform the user of relevant information, it responds to changes in the environment and, occasionally, it is proactive in that it may inform the user of information that is not exactly what had been asked for, but is judged to be sufficiently interesting to warrant informing the user.

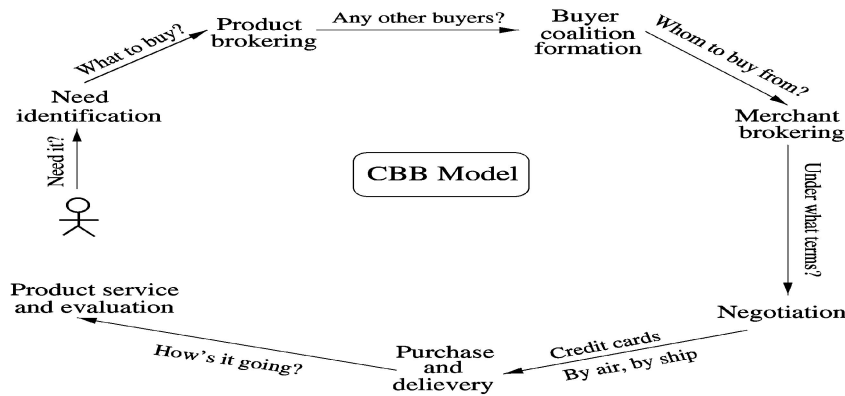


Fig. 2. Consumer behavior buying model.

TABLE 1
Filtering Techniques for Product Brokering in E-Commerce Systems

	Feature-based	Collaborative	Constraint-based
Amazon	✓		✓
eBay	✓		✓
CDNOW	✓	✓	
Yahoo shopping	✓		✓
Net Perceptions		✓	

TABLE 2
Comparisons of Different Product Brokering Techniques

	Feature-based	Collaborative	Constraint-based
When to use the technique	User's needs known	User's needs unknown	Some idea of user's needs
Requirements	Feature keywords for goods	Profiles of users	Conditions that goods satisfy
Interaction with user	Medium	Few	Medium
Results returned	Goods satisfying required features	Suggestions of goods to buy	Goods satisfying particular constraints
Goods suitable for	Most goods	Books, CDs, etc.	Most goods

shopper *A* is first compared with that of all the other shoppers in the system. Then, the “nearest neighbor” of *A* (i.e., the shopper whose profile is closest to that of *A*) is identified. Since shoppers with similar tastes and preferences are likely to buy similar products, the profile of the identified shopper is used to pass recommendations onto *A*’s agent. For example, in Net Perceptions (<http://www.netperceptions.com>), users are recommended the documents that their “knowledge neighbors” find valuable. In CDNOW (<http://www.cdnnow.com>), users are notified about the CDs or movies that are popular with other users with similar preferences. Constraint-based filtering involves an agent specifying constraints (e.g., the price range and date limit) to narrow down the products. In this way, customers’ agents are guided through a large feature space of the product [54]. For example, eBay guides a user agent to select the products by narrowing down the range of the possibilities based on the constraints the user gives (e.g.,

price range, item location, and so on). In the end, a list of the desired products that satisfy the user’s constraints is returned. Some e-commerce systems use more than one kind of filtering technique (since sometimes users do not know exactly the constraints of the goods they are looking for in advance). For example, eBay and Yahoo Shopping (<http://shopping.yahoo.com>) use both feature-based and constraint-based techniques. The differences among these techniques are summarized in Table 2.⁶

6. Most of the dimensions in Table 2 are self explanatory. However, for “interaction with user,” few interactions are needed in collaborative filtering since what the user agents need to do is just provide their user’s profile and they can then get recommendations from the system. For feature-based and constraint-based systems, some keywords or constraints need to be input until the user can find the exact product they want. The last dimension in the table is “goods suitable for.” Collaborative filtering is more specialized than the other techniques because it works based on perceived quality and people’s tastes rather than objective properties [131]. Thus, it is more suited to goods such as novels, CDs, and DVDs because it is subjective judgements that act as the differentiator in these cases.

2.3 Buyer Coalition Formation

Having determined the product to buy, customers may move directly to the merchant brokering phase (see below) or they may interact with other similar buyers to try and form a coalition before moving to the merchant brokering phase. Here, a coalition is viewed as a group of agents cooperating with each other in order to achieve a common task [134]. In these “buyer coalitions,” each buyer is represented by their own agent and together these agents try and form a grouping in order to approach the merchant with a larger order (in order to obtain leverage by buying in bulk). In [165], for example, a buyer coalition formation scheme is proposed in which buyer agents specify multiple items in a category and their valuation of these items and the group leader agent is then responsible for dividing the group into coalitions and calculating the surplus division among the buyers. Similarly, Tsvetovat and Sycara [149] view a buyer coalition model as being composed of five stages: negotiation, leader election, coalition formation, payment collection, and execution stages. They test their algorithms in a collective book purchasing setting in the university and show how the supplier agent gives a volume discount according to the size of the coalitions. In both of the above systems, it is essential to have a trustworthy and reliable agent that will collect the buyer’s information, divide the agents into coalitions, and negotiate with sellers (refer to [165], [149] for a full discussion of these issues).

2.4 Merchant Brokering

Having selected the desired product, and perhaps after having formed a buyer coalition, merchant brokering involves the agent finding an appropriate merchant to purchase the item from. Initial work in this area focused on finding the merchant that offered the good at the cheapest price. BargainFinder [80] was the first system of this kind to employ agents and it operated in the following way: If a customer wants to buy a music CD, BargainFinder will launch its agent to collect the prices from a predefined set of CD shops, and then it will select the CD with the lowest price for the customer. Another similar example is Priceline (<http://www.priceline.com>), which carries out the same set of tasks for airline tickets, hotel rooms, and cars.

However, in many cases, price is not the only determinant for the user. Other relevant issues, for example, might include delivery time, warranty, and gift services. Also, many merchants prefer their offerings not be judged on price alone. Thus, there is a move to extend these agents to consider multiple attributes. Naturally, the importance of the different attributes will vary between consumers and so there needs to be a way for this information to be easily conveyed to the agent. In the Frictionless Sourcing (<http://www.frictionless.com>) platform, “Vendor Scorecards” (multiattribute comparisons) are used to measure the performance of suppliers. For example, when evaluating the performance of different laptop computer suppliers, the key factors considered include reliability, responsiveness (e.g., reacting quickly), environmental friendliness (e.g., minimal pollution of the environment), and business efficiency (e.g., support for electronic purchasing over Internet). A total score is then calculated for each supplier based on the weighted score of these individual constituent components. These weights are obtained by the customer identifying themselves with a particular stereotype profile in which the weights are given.

2.5 Negotiation

Having selected a merchant (or set of merchants), the next step is to negotiate the terms and conditions under which the desired product will be delivered. We believe that one of the major changes that will be brought about by agent-mediated e-commerce is that dynamic pricing and personalization of offers will become the norm for many goods and customers. Thus, negotiation capabilities are essential for e-commerce systems [8]. In human negotiations, two or more parties bargain with one another to determine the price or other transaction terms [37]. In an automated negotiation, software agents engage in broadly similar processes to achieve the same end [65]. In more detail, the agents prepare bids for and evaluate offers on behalf of the parties they represent with the aim of obtaining the maximum benefit for their users. They do so according to some negotiation strategy. Such strategies are determined by the negotiation protocol that is in place. This protocol defines the “rules of encounter” between the agents [115], i.e., who can say what, to whom, and at what time. Given the wide variety of possibilities (as will be shown below), there is no universally best approach or technique for automated negotiations [65], rather protocols and strategies need to be set according to the prevailing situation [41]. Given this, our analysis of automated negotiation models as used in B2C e-commerce is divided into two categories:⁷ auctions (Section 2.5.1) and bilateral negotiations (Section 2.5.2).

2.5.1 Auctions

Auctions are one of the most widely studied and employed negotiation methods in e-commerce today [108]. For example, in the Internet Auction List,⁸ there are currently more than 2,500 auction company listings (2002). Online auctions make the physical limitations of traditional auctions disappear (e.g., time, space, and presence), and they provide millions of globally dispersed customers with more varieties of goods that can be selected within a flexible pricing mechanism [6]. Generally speaking, auctions are a very efficient and effective method of allocating goods/services, in dynamic situations, to the entities that value them most highly [163]. In an auction, one or more agents (called auctioneers) initiate the auction, and several (other) agents (called bidders) make bids according to the imposed protocol (which may permit one or multiple rounds). The outcome of the auction is then usually a deal between the auctioneer and the successful bidder.

There are many different types of auction (indeed, Wurman et al. [164] define a taxonomy of auction parameters that allows for approximately 25 million types of auctions). However, there are four common types of

7. There are other types of negotiation protocol such as multilateral negotiation (in which the negotiation involves bargaining between multiple noncooperative parties [1]) and n-bilateral negotiations (in which the negotiation involves multiple bilateral bargaining encounters [33]). However, since these protocols are not as widely used in e-commerce, we do not consider them here.

8. <http://www.internetauctionlist.com>.

TABLE 3
Values of Each Dimension in Table 4

Auction mode	one-sided (O) two-sided (T)	only bids or asks are permitted both bids and asks are permitted
Duration time	single-round (S) multi-round (M)	the auction only lasts one round the auction lasts multiple rounds
Unit of goods	one (O) many (M)	only one good is auctioned during the auction multiple goods are auctioned
Ratio of B-S	many to one (MO) one to many (OM) many to many (MM)	there are multiple buyers and only one seller there are multiple sellers and only one buyer there are multiple buyers and multiple sellers
Information revealed	yes (Y) no (N)	there is intermediate information revealed during the auction a bidder has no information about others
Settlement price	first price (F) second price (S) different prices (D)	highest price among all the bidders second highest price among all the bidders trades take place any time at different prices
Closing rules	time (T) inactivity (I) budget (B)	when a time is reached when there are no more bids for a time period when a reserve price is reached

single sided auction:⁹ English (first-price ascending), First-price sealed-bid, Vickrey (second-price sealed-bid), and Dutch (first-price descending) [122], [153]. The most common type of double-sided auction is the continuous double auction (CDA) (which allows buyers and sellers to continuously update their bids/asks at any time in the trading period).

The role of agents in such online auctions is to represent their users. This may involve monitoring the auctions, analyzing the market situations, and/or deciding when and how much to bid for the desired items. Automating these activities can save time, and in complex settings, we believe agents are likely to be more effective than human bidders. For example, Das et al. [27] show that agents outperform their human counterparts in a particular auction setting. Furthermore, a new possibility afforded by an agent-based approach is that a user can compete in multiple auctions simultaneously. Such a strategy has several advantages over participating in single auctions; for example, it can increase the chance of getting the good for customers, bring greater profit to customers by comparing multiple auctions and transacting at the cheapest price, and make the auction markets themselves more efficient by ensuring the transaction price is close to the equilibrium price [111].

Some good examples of auction websites include eBay, AuctionBot,¹⁰ and Yahoo Auctions,¹¹ and more information on auctions can be found in [9], [95], [98], [99], [160], [153]. To summarize, the characteristics of the auction types discussed above are listed in Table 4. The table consists of the seven dimensions explained in Table 3.¹² Given these

9. In single-sided auctions, only buyers or sellers can submit bids or asks; however, in double-sided auctions, buyers can submit bids and sellers can submit asks in the same marketplace.

10. <http://www.auctionbot.com>.

11. <http://auctions.shopping.yahoo.com>

12. Here, we only consider the popular forms of the auctions on same type, nondivisible goods, and the quantity is a single unit (see [163] for a fuller description of the bidding rules). Also, many different types of information can be revealed in the course of the auction (e.g., the identity of the bidders, the settlement price, the ask-bid spread, and so on [41]) here, however, we only consider whether any intermediate information is revealed (see [164] for a detailed discussion of the impact of information revelation).

protocols, we now turn to the strategies the agents need to employ in order to be successful.

- *English auction.* The agent's dominant strategy (the best thing to do, irrespective of what the others do [122]) is to bid a small amount more than the current highest bid and stop when the user's valuation is reached. For example, in Yahoo auctions, "automatic bidding" allows buyers to input their maximum bid (i.e., valuation of the item) and an agent will bid incrementally when it is necessary to win the auction.
- *First-price sealed-bid auction (FPSB).* In general, there is no dominant bidding strategy in this auction. Here, the price of the bid and the time to stop are functions of the agent's own valuation of the item and its beliefs about the valuation of other bidders. A good strategy is to bid less than the user's true valuation, but how much less depends on the user's attitude toward risk, the user's private valuation, as well as the prior beliefs about the valuations of other bidders. An analysis of such strategies can be found in [95].
- *Vickrey auction.* In a private value¹³ Vickrey auction, the dominant strategy is to bid the user's true valuation [122]. In this context, agents truthfully reveal their preferences which allows efficient decisions to be made.
- *Dutch auction.* The Dutch auction is strategically equivalent to the first-price sealed-bid auction. This is because in both games, an agent's bid matters only if it is the highest, and no relevant information is revealed during the auction process [122]. Klik-Klok (<http://www.klik-klok.com>) is an example of a Dutch auction Website for gold and jewellery sale where auction prices decline until a buyer makes a bid. The analysis of strategies in Dutch auctions can be found in [98].

13. Private value means that the bidder only knows its own valuation of the item, but it does not know the value to other bidders [95]. For example, auctioning a cake (that the winner will eat) is a private value auction, whereas auctioning treasury bills (where the value of the bills comes from reselling) is a public value auction.

TABLE 4
Comparison of Different Types of Auctions

Auction	auction mode		duration time		unit goods		ratio of B-S			information revealed		settlement price	closing rules
	O	T	S	M	O	M	MO	OM	MM	Y	N		
English	✓			✓	✓		✓	✓			✓	F	I
FPSB	✓		✓		✓		✓	✓			✓	F	T
Vickrey	✓		✓		✓		✓	✓			✓	S	T
Dutch	✓			✓	✓		✓	✓				F	B
CDA		✓		✓		✓			✓		✓	D	I

- Continuous double auction (CDA).* This type of auction is easy to operate, efficient, and can quickly respond to changing market conditions. A variety of different CDA models have been constructed [31], [46], [118], and these vary in terms of whether bids/asks are for multiple or single units, whether unaccepted offers are queued or replaced by better offers, and so on [41]. Nevertheless, all these protocols allow traders to make offers to buy or sell and to accept other traders' offers at any moment during a trading period [41]. The messages exchanged generally consist of bids (offers to buy) and asks (offers to sell) for single units of the commodity, and acceptances of the current best bid or ask. Several bidding strategies have been proposed in the literature. The ZERO Intelligence strategy [50] generates a random bid within the allowed price range decided by the agent's budget constraint. The adaptive agent bidding strategy is based on stochastic modelling of the auction process using a Markov chain [106], [107]. A sequential bidding method that uses dynamic programming is proposed in [146]. The steal strategy involves "waiting in the background and bidding when the bid and ask get very close" [117]. In [58], heuristic fuzzy rules and fuzzy reasoning mechanisms are used to determine the best bid to make given the state of the marketplace.
- Multiple auctions.* In this type of auction, the agent needs to monitor all the relevant auctions, decide which one to bid in, and determine what to bid in order to get the goods at the best deal. There are a variety of tools available for tracking the various auctions that occur for the desired item. For example, BidXS (<http://www.bidxs.com>) can search more than 500 auction sites for the item the user wants. Again, there is no dominant bidding strategy, thus heuristic methods are required. Preist et al. proposed a coordination algorithm designed for agents that participate in multiple simultaneous English auctions [111], [110]. BiddingBot is a multiagent system that supports users in attending, monitoring, and bidding in multiple auctions through cooperative bidding [62]. Anthony and Jennings [3] present the design of a heuristic agent that participates across multiple English, Dutch, and Vickrey auctions. In [19], a framework that enables

an agent to make rational decisions across multiple simultaneous auctions is developed.

2.5.2 Bilateral Negotiations

Bilateral negotiation involves two parties, a service/good supplier and a consumer, coming to a mutually acceptable agreement over the terms and conditions of a trade [136]. In contrast to most of the auction work (which is a form of one to many, many to many, or many to one negotiation), bilateral negotiation is usually concerned with multiattribute contracts (covering price, quality, delivery date, and so on). As with the auction work, there is no dominant negotiation model or strategy that is suitable for all occasions. Rather, it is a case of different models having different strategies that are suitable in different contexts. Given this situation, we classify extant work on bilateral negotiation into three groups.

- Decision making by explicitly reasoning about the opponent's behavior.* Agents using the strategies in this group explicitly reason about their opponent's objectives and behaviors and then decide what is the appropriate response to their likely behavior. In this respect, noncooperative game theory (which is particularly concerned with providing equilibrium strategies in which no agent wants to change its strategy whatever its opponents do) is an important approach for analyzing strategic interactions among agents [79], [148]. The recursive modeling method [154], [49] is employed by an agent to reason about its opponent so that it can generate its own strategy in response. In [171], a Bayesian network is used to update the knowledge and belief that each agent has about the environment and other agents; and offers and counteroffers between agents are generated based on Bayesian probabilities.
- Decision making by finding the current best solution.* Algorithms in this group focus on finding the offer/counteroffer that maximizes the agent's profit given the agent's constraints, preferences, current negotiation situation, and the opponent's last offer. In Tête-à-Tête [53], constraints on product features and constraints on merchant features are used to influence the decision of what and whom to buy from. Luo et. al [89] developed a fuzzy constraint-based framework for multiissue negotiations in competitive trading environments and demonstrated it in a negotiation between a real estate agency and a customer. Kowalczyk and Bui [77], [76] also use fuzzy constraints to model

TABLE 5
Dimensions of B2C and B2B E-Commerce

Dimension	Business-to-customer e-commerce
Features	<ul style="list-style-type: none"> • Individual customer oriented • No collaboration between customer and company required • Brings convenience for buying in globally competitive markets • Quick response to the transaction • Convenient to use
Roles of agents	<input type="checkbox"/> Need identification <ul style="list-style-type: none"> • <i>e.g.</i>, Amazon (Delivers) and Fastparts (Auto Watch)
	<input type="checkbox"/> Product brokering <ul style="list-style-type: none"> • Feature-based, collaborative and constraint-based filtering • <i>e.g.</i>, Amazon, eBay, and Net Perceptions
	<input type="checkbox"/> Buyer coalition formation <ul style="list-style-type: none"> • <i>e.g.</i>, Collective book purchasing and GroupBuyAuction
	<input type="checkbox"/> Merchant brokering <ul style="list-style-type: none"> • Price comparison and multi-attribute comparison • <i>e.g.</i>, BargainFinder and Frictionless Commerce
	<input type="checkbox"/> Negotiation <ul style="list-style-type: none"> • Auctions and multi-attribute bilateral negotiation • <i>e.g.</i>, eBay, Yahoo Auctions and AuctionBot
Dimension	Business-to-business e-commerce
Features	<ul style="list-style-type: none"> • Organisation oriented • Close collaboration between organisations required • Facilitates both direct and indirect procurement and supply chain • Larger and global markets • Real-time and low cost transaction • Less inventory and dynamic pricing
Roles of agents	<input type="checkbox"/> Partnership formation <ul style="list-style-type: none"> • Virtual enterprises • Supply chain management
	<input type="checkbox"/> Brokering <ul style="list-style-type: none"> • Information retrieval & processing, negotiation, Profiling of users, notification, collaboration with other brokers • <i>e.g.</i>, OFFER, MULTIMEDIATOR, and Abrose
	<input type="checkbox"/> Negotiation <ul style="list-style-type: none"> • Auctions (sell-side, buy-side and combinatorial auction) • <i>e.g.</i>, Fastparts.com, GE.com, Ariba.com and labx.com • Contracting • <i>e.g.</i>, Contract net, marginal cost-based contract, OCSM levelled commitment contract and MAGNET

multiissue negotiation, but their approach performs negotiation on individual solutions one at a time. Matos et al. [94] employ fuzzy logic, case-based reasoning, and evolutionary computing to deal with the bilateral negotiation [66]. Faratin et al. develop a suite of algorithms for multiissue negotiation that covers both concessionary behavior [34] and trade offs that aim to find a win-win solution for both parties [35].

- *Argumentation.* In the argumentation-based approach, agents exchange additional information over and above the basic terms and conditions of the contract [65]. This information can be of a number of different forms, nevertheless, it is always some form of argument which explains/justifies the position of the agent making the argument. Thus, in addition to rejecting a proposal, an agent can offer a critique of the proposal, explaining why it is unacceptable (*e.g.*, the price is too high). The way in which argumentation fits into the general negotiation process was defined in [137] where a simple negotiation protocol

for trading proposals was augmented with a series of illocutionary moves which allow for the passing of arguments.

3 AGENTS IN B2B E-COMMERCE

Compared with B2C e-commerce, B2B deals with transactions among organizations (see Table 5 for a more detailed comparison). Generally speaking, relationships between organizations are more complex than those between businesses and consumers since they involve the adoption of similar standards with respect to communications and collaboration, as well as joint information technology investment [141]. In particular, one of the main aims of B2B e-commerce is to significantly improve the supply chain by facilitating more efficient and agile procurement processes [30]. Moreover, the exchanges in the B2B domain are increasingly tending to be private [167]. Such exchanges enable companies to trade with their existing partners in a well-defined environment without having to go through some of the early stages of the B2B lifecycle (see Fig. 3).

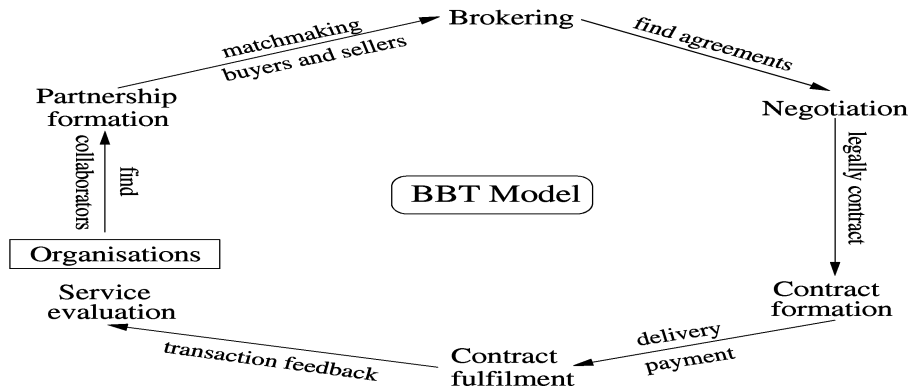


Fig. 3. Business-to-business lifecycle model.

There are various ways of viewing B2B e-commerce (e.g., [152], [7]), however, here we use a BBT model (see Fig. 3) to explore the roles of agents. Specifically, we believe agents are most useful in the partnership formation, brokering, and negotiation stages because these stages all involve complex issues related to decision making, searching, and matchmaking that agents are well suited to. Thus, we will explore these roles in more detail in the rest of this section. Currently, agents are not used in the contract formation stage, however, we believe they have the potential to be involved in this activity. Contract formation marks the termination of negotiation and involves the agreed terms being put into a legally binding contract. Traditionally, contract formation involves two or more people, meeting face-to-face. However, as e-commerce systems evolve, this situation is starting to change. In the US, for example, Section 206 in the Uniform Commercial Code (U.C.C.) was proposed by reformers as a way of dealing with automated contract formation and clearly states that contracts can be formed by the interaction of electronic agents.¹⁴ This proposal is motivated by the fact that it is a challenging task for the courts to determine where the communication system ends and when the legal agent begins.¹⁵ The fifth stage, contract fulfillment, means the parties carry out the agreed transaction according to the terms specified in the contract. This stage usually includes: a detailed description of the good/service provided; the means of delivery (electronic or physical); how it will be paid for (e.g., partial payments up-front, with the balance paid on completion); which law governs the contract; how to resolve any disputes, how to deal with claims arising, how a contract can be monitored, and so on. We believe agents are not likely to be involved in this stage for some time because it involves many complex legal issues and subjective judgments. The last stage, service evaluation, is the posttransaction stage, where traders evaluate their satisfaction with the transaction. Many e-commerce systems allow users to provide feedback on the transactions experienced. For example, eBay uses "Feedback Forum" to check the reputation or business practices of anyone at eBay. This

feedback, representing the reputation of the trader, can then be made accessible to subsequent agents that wish to interact with the trader. Again, because of its subjective nature, we do not believe there is a significant role for agents in this phase of the lifecycle.

3.1 Partnership Formation

The information technology available today makes it possible for a company to search for its partners worldwide [81]. Given this fact, partnerships can be much more agile and fluid. Thus, this step may include the forming of a new virtual organization as well as finding the partners that provide products or services in a supply chain.

3.1.1 Virtual Enterprises

A virtual enterprise¹⁶ (VE) is composed of a number of cooperating companies that share their resources and skills to support a particular product or project effort (for as long as it is viable to do so) [105]. The idea is that by collaborating, the constituent companies can more effectively utilize their resources than if they acted in isolation [51]. For example, an individual company may collaborate with several partner companies that provide related products so that each of them need only provide the services/products in which they specialize, but, when taken together, the VE can provide a broader range of offerings. Such VEs offer several potential advantages [92]: maximizing flexibility and adaptability to respond to environmental changes; developing a pool of competencies and resources by combining its members' resources; adjusting itself according to the market constraints; and managing the global supply chain optimally.

Given the fact that a VE is composed of a number of autonomous entities that need to interact with one another in flexible ways, agent technology is a natural underpinning model [105]. In more detail, the formation of a VE involves a selection process based on a number of variables such as organizational fit, technological capabilities, relationship development, quality, price, and speed [125]. Thus, a broker may assist in identifying the best partners from the set of potential collaborators [97] (see also Section 6). Having

14. The official draft of Article 2B of the U.C.C. is from 2002 and can be found at <http://www.law.uh.edu/ucc2b>.

15. See <http://www.jurisdiction.com/ecom3.htm> for more details of this debate.

16. Coalition formation by buyer companies is similar to buyer coalition in the B2C domain (Section 2.3). Thus, here we focus on virtual enterprises of supplier agents.

identified the partners, the agents need to negotiate with one another in order to set the terms and conditions of their partnership [151] (and Section 3.3). Then, once the VE is established, the agents need to coordinate their actions so that they deliver their services in an effective manner. Here, the VE might require a number of agents to manage its ongoing operations [93]. For example, Martinez et al. [92] propose a multiagent control system that consists of three kinds of controller agent: product agents (which manage the activity associated with each product), activity agents (which autonomously manage an entire manufacturing activity), and resource agents (which manage their own operative functions and propose service offers to activity agents). Together, these agents use and control the other entities in the system in order to achieve the VE's overall aims. The MASSYVE (multiagent agile manufacturing scheduling systems for virtual enterprises) project focuses on the use of multiagent systems in agile scheduling in a VE environment [112]. The factors considered here range from distribution logistics scheduling in supply chains to negotiation in the VE using mobile agents. The AIMS (agile infrastructure for manufacturing systems) project enables companies to share resources and skills to facilitate the operations of VEs and agents function as a bridge between clients and servers [105]. Specifically, the agents act as: facilitators (routing requests to appropriate databases), aggregator agents (combining multiple orders), user programmable agents (automating routine tasks), and engineering databases agents (notifying users of design changes).

3.1.2 Supply Chain Management

A supply chain is formed by business units or facilities that purchase raw materials, convert them into intermediate goods and final products, and delivers these final products to customers [143]. A supply chain is used to coordinate the activities of the organizations involved in order to ensure that products pass through the chain in the shortest time and at the lowest cost [84]. Because of the business trend toward outsourcing services and resources, supply chain networks have become more complex [81]. Given this, the software solutions being developed need to be more sophisticated than the current generation of workflow tools. In particular, the various components of the supply chain can be viewed as autonomous stakeholders and these various stakeholders need to interact in flexible ways. Thus, an agent-based approach is well suited to this domain [60], [66]. In particular, agents can be used to execute the scheduling [40], negotiate about product prices [142], and share data between companies [170].

To this end, a number of models for agent-based supply chain management have been reported. For example, Walsh and Wellman developed a market system, based on a task dependency network, for allocating tasks among agents that compete for scarce resources [157]. Sun et al. [142] model and implement the order selection and negotiation process in a supply chain as a multiagent system. Here, a negotiating agent represents each company along the supply chain and agents generate a purchase plan, negotiate, and generate counter proposals using constraint satisfaction techniques. Moreover, Zeng proposed the

Leadtime-Cost Tradeoff Supply Chain Model [170] in which each agent represents a business entity and the agents coordinate with each other to control activity in the supply chain. An agent-based approach for streamlining the business decision process is proposed in [73]; here, agents assist the decision maker by discovering matches between supply and demand.

3.2 Brokering

Brokering is the process that matches sellers who supply goods/services to the buyers who need them [38]. From the seller's side, it is how they can propagate their products and locate potential buyers. From the buyer's side, the problem is how to find the most appropriate seller to provide the good/services [74] (e.g., lowest price or best quality). In contrast to merchant brokering in the B2C domain (Section 2.4), brokering in a B2B context typically involves repeated transactions and large volumes (in the B2C context, brokering requests often tend to be one-off transactions since individual customers tend not to buy the same product often).

As the Internet expands, it is becoming more expensive and more difficult to navigate in order to find the necessary information on companies and their offerings (this is especially true for small and medium size enterprises [152]). For example, it is estimated that about \$5 trillion dollars is spent on the procurement of industrial parts each year [150]. Given the difficulty and value of this exercise, a common way of obtaining this information for companies in B2B e-commerce is through some form of information broker (they can also be called matchmakers [55] or brokerage centres [44]) that acts as an intermediary between the buyers and sellers. Here, a broker can be an agent or a multiagent system. The functions offered by a broker may include the following [38]: information retrieval and processing, maintenance of a self-learning information repository about the user, profiling of users, monitoring for items of interest to the user, filtering of information, intelligent prediction of user requirements, commercial negotiation between customers and the providers, collaboration with other brokers, and protecting the user from intrusive access.

To summarize, Table 6 details the services provided by a number of agent-based broker systems. As shown in the table, most brokerage services today mainly focus on information search and matchmaking buyer's and seller's profiles, as well as comparing the products in the catalogs of different suppliers. We believe more advanced services (e.g., collaboration with other brokers and protecting the user from intrusive access) will now start to emerge in order to provide more support to the buyers and sellers involved in the transactions.

3.3 Negotiation

After the appropriate providers and consumers have been brokered, the negotiation stage is where the traders aim to reach an agreement about what actions should be performed under what conditions. By establishing contracts on an as-needed, just-in-time basis, sellers can tailor their offerings both to their individual and the prevailing market situation at any given moment in time. Buyers can reduce

TABLE 6
Brokers' Functions and Services

Project/ Papers	Information retrieval and processing	Maintenance self-learning about user	Profiling of users	Nego- tiation	Collaboration with other brokers	Notifi- cation
Abrose [47]	✓	✓	✓		✓	
OFFER [14]	✓		✓	✓		✓
MULTIMEDIATOR [43]	✓	✓	✓			
Schmid <i>et al.</i> [128]	✓			✓		

their supply chain cost, benefit from dynamic pricing mechanisms, broaden their supplier database, and streamline the procurement process. Compared with negotiation in the B2C context, B2B negotiation is more complex. Typically, for example, it involves larger volumes, repeated transactions, and more complex contracts. The negotiation methods discussed in the B2C context can also apply here, however, the two most popular means of conducting B2B negotiation are through auctions (Section 3.3.1) and contracting (Section 3.3.2).

3.3.1 Auctions in B2B E-Commerce

There are now many B2B marketplaces on the Web that provide auction services and allow organizations to trade with one another on a global basis (e.g., FreeMarkets¹⁷ and Ariba¹⁸). Indeed, industry analysts estimate that 25 percent of e-commerce now consists of exchanges through such mechanisms [126]. These auctions offer many advances over traditional exchange methods (e.g., fixed suppliers), such as a larger market, less inventory, reduced transaction costs, global expansion, and efficient pricing [126]. We classify the commonly used auctions into three kinds: buy-side auctions (one buyer and multiple sellers), sell-side auctions (one seller and multiple buyers), and combinatorial auctions [42], [72], [113] (where bidders bid for a combination of related items). An agent can be either a buyer who submits bids or a seller who provides some products or services in these auctions. The sell-side auction is similar to the auctions discussed in the B2C context; the buy-side auction is the opposite of the sell-side auction (however, it can also be an English, Vickrey, FPSB, or Dutch auction). Combinatorial auctions only take place in B2B environments because of their inherent complexities.

Buy-side Auctions. Buy-side auctions, also called reverse auctions [144] or procurement auctions [24], occur when buyers negotiate with multiple sellers in order to procure a particular good/service.¹⁹ In this case, the negotiation usually involves multiple attributes since buyers invariably have their particular requirements on the goods/services they need. Here, the buyer sends out his

requirements and the sellers who can meet them make bids. To make this process cost effective, some companies have built their own markets in which they can invite bids from potential sellers. Examples of this kind include General Electric (<http://www.ge.com>) and Boeing Inc. (<http://www.boeing.com>). The idea is that the cost spent in searching and comparing suppliers can be significantly reduced because the companies repeatedly buy large volumes of many such products. In contrast, some companies conduct buy-side auctions through a third party Website (e.g., labx and Ariba).

More specifically, Che [24] investigates government procurement using a two-dimensional auction (price and quality). A buyer solicits bids from multiple sellers. Each bidder submits a sealed bid specifying the price and quality and the bidder with the highest score wins. Based on the different ways in which the winner offers the goods/services, three auction schemes are proposed: first score (winner offers the price and quality it bids), second score (winner offers the goods/services matching the score of the second highest scored bidder), and second preferred offer (winner offers the goods/services at the same price and quality as the second highest scored bidder). In this model, the buyer evaluates the bids by a scoring function which converts a bid into a single number.

Bichler et al. [13] also defines a bidding procedure for multiattribute auctions: a buyer first specifies a request for bids and defines his requirements and preferences for the goods/services in a scoring function. Then, sellers submit their bids. After the auction closes, the winning bid is the one that has the highest score as computed by the buyer (i.e., the seller who satisfies the buyer the most). This basic mechanism was applied to English, Vickrey, and FPSB auction protocols. Moreover, Bichler et al. empirically analyzed these multiattribute auctions and found that the utility scores achieved by the buyer are significantly higher than those of the corresponding single-attribute auctions [11]. In this setting, the scoring function is revealed to the bidders, thus the bidders know how to improve their bids in a way that makes them most attractive to the buyer, and least costly for them.

Teich et al. [144] developed multiple-attribute algorithms and heuristics for auctions. In the case where the "quantity" is not an attribute in the auction, the preference of the auctioneer is represented by the preference path that is the ordering of all the levels of each attribute. The preference path here acts as a scoring function that the bidders would follow. In the case where "quantity" is an attribute in the auction, a discriminative auction algorithm is proposed.

17. <http://www.freemarkets.com>.

18. <http://www.ariba.com>.

19. There is another form of buy side auction called a Request for Quote (RFQ) [152]. In an RFQ, the buyer requests quotes that can include the price, delivery dates, and description of the goods or services being provided. The buyer uses this as a way to begin negotiation. However, since there is not an automatic criteria (e.g., a scoring function to evaluate the bids) for selecting the winner, the strategy for the bidder is not obvious. Thus, we only discuss the reverse auction in this context because it has a clear selection criteria which means it is amenable to an agent-based solution.

The auctioneer can specify multiple reservation prices for different quantities and the bidders can accept the suggested bid or bid above the suggested bid price. The authors argue that the algorithm can make the market more efficient.

Vulkan and Jennings [156] proposed a multiattribute auction protocol for service allocation in the ADEPT [66] scenario. The ADEPT technology was used to develop a system for managing the British Telecom (BT) business process of providing a quotation for designing a network to provide particular network services to a customer. They show that the protocol is guaranteed to choose the service provider that makes the best offer from the buyer's utility perspective and that this offer is better than any offer that would have been forthcoming using any other negotiation protocol.

Sell-Side Auctions. In this kind of auction, there is a seller who wants to sell goods/services and many buyers join the auction. The mechanism is usually one of the common single sided types described in Section 2.5.1 and the strategies described in that section also apply here. These sorts of auctions are often used by companies that hold excess inventory or that buy out-of-date inventory [126]. Fastparts, which sells electronic manufacturing products, and Staples (<http://www.staples.com>), which sells business supplies and services, are two of the most prominent examples of this genre.

Combinatorial Auctions. These are a special form of auction in which there are multiple kinds of goods to sell and bidders can bid on combinations of items. For example, a seller may want to sell several kinds of related goods (e.g., licenses in spectrum auctions²⁰) and many bidders may have preferences over a combination of items (e.g., bidding on license A and B for \$300). After the seller receives all the bids, it will decide a nonconflicting allocation among these goods that maximizes its revenue. These sorts of auctions are involved in many situations in the real world. For example, in the Federal Communications Commission (FCC) spectrum auction,²¹ bidders placed bids on different combinations of spectrum licences. Between 1994 and January 2002, 38,829 licenses have been auctioned and 21,849 of them have been won through such combinatorial mechanisms. Other examples of combinatorial auctions are for airport time slots [113], railroad segments [18], and delivery routes [20].

These auctions are especially prevalent in a B2B context because companies often want to trade in a variety of interrelated assets. Moreover, as different companies value the items or bundles of items differently, allowing them to bid on combinations provides greater flexibility in expressing their needs and enhances the economic efficiency of the market [155]. From an agent perspective, the key challenge is that of winner determination (auctioneer selects a set of nonconflicting bids that maximize its revenue) and a number of algorithms have been developed to achieve this according to various criteria (e.g., anytime algorithms [42], polynomial algorithm [26], and optimal solutions [124]). In this context, the bidding agent has to express its preferences

on every bundle it is interested in. However, transmitting these preferences to the auctioneer is a difficult task since the bundles in the bids are likely to be very large. To overcome this, some researchers have developed an "oracle" (a program that can compute the bid for each bidder) [155] and others have developed a bidding language to encode the preferences of the bidders (e.g., XOR-bids and OR-of-XORs [124]).

3.3.2 Contracting

Contracting covers the negotiation involved in reallocating work among agents; it involves one agent trying to contract out some of its tasks to another agent by promising some rewards [78]. Contracts have been applied in fields such as electricity markets, bandwidth allocation, manufacturing planning and scheduling, and electronic trading of financial instruments [121].

Smith's Contract Net Protocol [139] was the first multiagent contracting protocol. In this protocol, a manager agent announces a task, receives and evaluates bids from potential contractors, and then awards the task to one of them and finally receives the result from this contractor. Sandholm extends this work to consider marginal cost-based contracts (an agent contracts in/out a task only if it can make a profit doing so) [119]. Sandholm's protocol was used as the basis of the TRACONET system which is an automated system for task reallocation among freight companies. Here, each agent, representing a company, can take delivery tasks from or give out tasks to other agents. In the original contract, only one task can be moved between agents at any one time. However, this sometimes led to local optima. To overcome this, several new types of contract were added [119]: cluster contracts (exchanging multiple tasks), swap contracts (swapping a task for another task), and multiagent contracts (more than two agents in the same contract). When taken together, contracts that combine all of the above can be shown to guarantee the optimal allocation through a finite number of contracts. Andersson and Sandholm [2] also devised leveled commitment contracting in which an agent can decommit from contracts by means of paying a monetary penalty to the contracting partner as a way of releasing itself from the contract. Collins et al. [25] developed the MAGNET (Multi AGent NEgotiation Testbed) system which takes advantage of an independent market infrastructure and uses it as an intermediary to facilitate the interactions between agents. Compared with other negotiation approaches (e.g., contract net), the fact that there is an explicit intermediary reduces counter speculation by enforcing negotiation rules and verifying the identity of the agents.

4 UNDERPINNING TECHNOLOGIES

This section briefly discusses some of the languages that are involved in supporting interactions among agents and some of the development tools and platforms for agent-mediated e-commerce. Finally, we discuss the trust issues involved with agents when selecting transaction partners. The aim here is not to present a complete overview of agent technology (such surveys can be found in [67], [103], [101]), but rather to highlight some of the key issues as they relate to agent-mediated e-commerce and to identify

20. <http://wireless.fcc.gov/auctions>.

21. <http://www.fcc.gov/wtb/auctions/>.

some of the technical solutions that currently exist to support it.

4.1 Interaction Languages and Protocols

An agent needs to be able to communicate with other agents, with the environment, and with humans. To realize this, a number of interaction languages have been developed. Many of the applications discussed use the eXtensible Markup Language (XML) [83] to code the information and services in meaningful structures that agents can easily understand and process. However, for XML to be widely accepted, the ontology problem (the fact that there needs to be an agreement on the meaning of the terms in the area of interest [61]) must be solved first, so that the tagged data in XML can be semantically consistent. Specifically in the e-commerce domain, the Electronic Business eXtensible Markup Language (ebXML)²² is based on international standards and aims to provide an XML-based open technical framework enabling data exchange for B2B and B2C e-commerce. Currently, however, most B2B transactions exchange information via the Electronic Data Interchange (EDI) protocol (which is complex and expensive [152]). Nevertheless, many people believe that the XML/EDI message format will eventually replace the traditional EDI messages [48]. Besides XML/EDI, there are a number of other XML-based specifications for domain-specific commerce languages. For example: Information Content and Exchange (ICE),²³ Open Buying on the Internet (OBI),²⁴ Secure Electronic Transactions (SET),²⁵ Open Trading Protocol (OPT),²⁶ and Open Financial Exchange (OFE).²⁷

The semantic Web provides a new solution for the development of Web-based intelligent agents. There are two semantic Web languages: Resource Description Framework (RDF)²⁸ and Darpa Agent Markup Language (DAML).²⁹ RDF Schema defines the ontology for RDF. Recently, a number of languages attempt to extend the expressive capability of RDF and RDF schema; these languages include SHOE,³⁰ DAML [59], and OIL [57]. Further, the best features of these languages are integrated in a hybrid to form the semantic markup language Web Ontology Language.³¹ More information about the ontology languages can be found in [52].

The above interaction languages are primarily defined for the Web. However, the agent field also has its own communication and knowledge representation languages. Interactions between agents (e.g., during the course of a negotiation) are often enacted via an agent communication language (ACL). ACLs are often based on the speech act theory [4], [129], which explicitly identifies the intention its speaker is trying to convey to the hearer. The most widely employed ACLs are the Knowledge Query Manipulation Language (KQML) [36] and FIPA ACL [104]. These two are broadly similar in their basic conception, but differ in some

of the primitives (e.g., there is no “facilitation primitives” like *broker* or *recommend* in FIPA ACL).

While ACLs are concerned with the format and intent of the message, the content of the message needs to be represented in some form of content language. Here, KIF (the Knowledge Interchange Format)³² and FIPA SL (Semantic Language), the content language for FIPA ACL, are the most commonly used languages. More information about agent communication languages can be found in [36], [82], [75], [158].

Moreover, there are a number of products and services on the Internet that aim to facilitate the cooperation among trading partners as well as promote the efficiency of the supply chain. These efforts include the Global Commerce Initiative (which aims to promote a global supply chain efficiency),³³ the UCCnet³⁴ (which enables the formation of collaborative relationships of trading partners), and the GlobalNetXchange (which helps retailers and suppliers improve efficiencies and reduce costs throughout the supply chain).³⁵

As this section has highlighted, there are many choices and options available for dealing with interactions in the agent-mediated domain. At this stage, it is not possible to determine which of these candidates is likely to become the dominant approach. However, with the increasing convergence of agent technology, the semantic Web [10], and grid computing [39], we believe that a small number of languages will soon become de facto standards for high-level interactions and knowledge exchange since such capabilities are at the heart of all these large-scale, open distributed systems [116].

4.2 Development Tools, Technologies, and Platforms

There are many general purpose agent development toolkits (see [54], [158] for more details), but here we focus on those that have been specifically developed for the e-commerce domain. Again, our aim is not to be exhaustive, but rather to give an indication of what is available. First, there are some agent-based negotiation servers. For example, AuctionBot³⁶ is a multipurpose Internet auction server that can be used to create automated auctions based on users' specifications. eMediator [123] is an e-commerce server that provides services to facilitate efficient commerce (e.g., eAuctionHouse (a third party auction site which provides a wide range of customizable auction types), eCommitter (leveled commitment contract (see Section 3.3.2) optimizer), and eExchangeHouse (a safe exchange planner)). FishMarket³⁷ is a multiagent test-bed for auction tournaments where agents can trade in an auction house using various auction protocols. Second, there are some agent development platforms and tools. These include Agent development kit (ADK),³⁸ a java-based e-commerce, data warehouse, and workflow management applications environment that consists of a set of the agent foundation classes and an agent runtime environment that together

22. <http://www.ebxml.org>.

23. <http://www.oasis-open.org/cover/ice.html>.

24. <http://www.openbuy.org>.

25. <http://www.setco.org>.

26. <http://www.otp.org>.

27. <http://www.ofx.net>.

28. <http://www.w3.org/RDF/>.

29. <http://www.daml.org/>.

30. <http://www.cs.umd.edu/projects/plus/SHOE/>.

31. <http://www.w3.org/TR/2002/WD-owl-guide-20021104/>.

32. <http://logic.stanford.edu/kif>.

33. <http://xml.coverpages.org/gci.html>.

34. <http://www.ucnet.org/>.

35. <http://www.gnx.com/home.jsp>.

36. <http://tac.eecs.umich.edu/auction>.

37. <http://www.iiia.csic.es/Projects/fishmarket/>.

38. <http://develop.tryllian.net/public/index.html>.

allows application programmers to define the components required to build an agent based application. E.piphany E.5³⁹ provides an environment for agent-based customer relationship management that coordinates, in real time, all inbound and outbound interactions with customers. It can be used in both the B2C and the B2B categories. Finally, some agent-based e-commerce solutions are also available. For example, the Lost Wax e-Commerce Platform (<http://www.lostwax.com>) supports buyers and sellers in both public and private trading environments and the agent-based modules can represent differing trading mechanisms (such as auctions, contracting, and negotiation). The living markets platform (<http://www.living-systems.com>) is an agent-based product for real-time optimization of processes in business networks.

In contrast to the position with interaction languages, we do not believe that one, or even a small set of, toolkits will become dominant in this sector. Rather, we believe that the requirements and functionality are sufficiently diverse that there will continue to be many offerings that are targeted at particular niches.

4.3 Trust Issues

In addition to the security issues that are common to all e-commerce applications [45], [96], the issue of trust becomes particularly important when we are considering the agent-mediated variety of e-commerce. This is because software agents have increased degrees of autonomy and because they engage in flexible interactions that are not necessarily foreseen at design time. In this context, the level of trust is determined by: the degree of initial success the agents experience with each other, the well-defined roles and procedures for all parties, and the realistic expectations as to the outcomes from the trading endeavor [152]. For instance, in the delivery process for transaction information, there should be some mechanism to ensure that both buyers and sellers keep to their side of the deal. Generally speaking, the solutions to trust problems fall into two classes. First, third-party enforcement [152]. In most agent-mediated e-commerce transactions, the issues of traceability and trust are handled by setting up trusted third parties such as public certification authorities (CA), banks, and credit-card companies. For example, agents can check the identities of the potential partners through the public CA before transactions. Second, unenforced e-commerce transactions. For example, Sandholm [120] presents a method in which the seller agent splits larger exchanges into smaller parts and the buyer agent pays only part of the money, and this exchange proceeds simultaneously so that no agent is motivated to defect at any point. Telser [145] has also analyzed safe exchanges of goods and payments in the similar way, however, he analyzed a sequence of exchanges as a repeated game.

The issue of trust is often mixed in with related issues such as reputation [100], source reliability, and deception (e.g., [16], [21], [127], [159], [161]). Since the only rational basis of agent design in this domain is self-interest, trust and the related issues are something that needs to be engendered in the agent being interacted with (e.g., the chosen rules and the guarantees offered by the third party).

This means that agents should be able to evaluate the risks of a collaboration, to select the best and most reliable partnership in a specific situation, and to revise previous unsuccessful collaborations in an adequate trust model [22]. There have been a number of both theoretical and empirical attempts to achieve this. For example, Zacharia and Maes [168] develop a collaborative reputation mechanism that offers personalized evaluations for the various ratings (assigned to users) to predict their reliabilities in an e-commerce context. Sen et al. [130] consider sharing other agents' "opinions" as a means of curbing the exploitative tendencies of selfish agents. To avoid the problems caused by believing other agents, they developed a learned trust-based evaluation function that can resist both individual and concerted deception on the part of selfish agents. Finally, Banerjee et al. [5] employ Bayesian networks to model the relationships among the agent dispositions and their actions. The agent updates its belief about other agents through observations in tandem.

5 CONCLUSIONS AND FUTURE CHALLENGES

This paper has surveyed and analyzed the state of the art in agent-mediated e-commerce, focusing particularly on the B2C and the B2B context. While agent-mediated e-commerce is still very much in its infancy, a number of agent-based deployments have already been made. In highlighting these endeavors, we have also tried to outline medium and longer term aspirations for this area. Specifically, however, there are a number of major research challenges that still need to be overcome before the full potential of agent-mediated e-commerce can be met.⁴⁰ These include, in no particular order:

1. **The personalization problem.** One of the key advantages of agent-mediated e-commerce is the personalization opportunities that the agent metaphor offers. Thus, rather than presenting the same options and making the same decisions for each and every user, the agent-based approach offers much greater tailoring to the individual user. Such personalization is often represented by an interface agent that acts as a personal assistant for the user in some task environment (see [29], [85], [87], [91] for more details of work in this area). However, for this to be effective, the agents need to know the following information about their user: their objectives, their preferences for particular goods, and the range of activities. However, techniques for acquiring this information in a nonobtrusive manner that does not overly burden the user are still the subject of ongoing research.
2. **The semantic interaction problem.** Getting agents that have not been predesigned to interwork to interact with one another in a meaningful and sophisticated manner is one of the major challenges

40. While there are already systems in which these issues have been overcome in some way, the solutions tend to be ad hoc, non-scalable, or limiting in some way. In raising these as research challenges, we are talking about developing robust and reliable solutions that are reasonably generic in nature.

39. <http://www.epiphany.com/products/index.htm>.

for computer science research. However, with the increasing convergence of agent technology, the semantic Web and grid computing considerable research impetus is building up behind this problem. This will, we believe, lead to widely applicable approaches and solutions in the medium-term.

3. **The discovery problem.** As the numbers of producers, consumers, and e-marketplaces increase, so it is becoming increasingly difficult to put the relevant agents in contact with one another. This is especially acute at the product brokering (Section 2.2) and merchant brokering (Section 2.4) stages. This problem of scale is then further exacerbated by the dynamic nature of agent-mediated e-commerce in which sellers and buyers can freely enter and leave the system.
4. **The interaction problem.** Many of the interaction models that agents employ are still reasonably simplistic in nature (especially in the deployed systems). More research is needed to develop better models for the more complex forms of auctions and for the bilateral negotiation scenarios. As well as producing more efficient models, further work is also needed to give better predictive capabilities about how a given model will perform in various types of environments.
5. **The trust problem.** Trust has both a social and a technological facet and both of these need to be further addressed if users are to be happy to delegate increased autonomy to a software agent that interacts on their behalf. From a social perspective, people will need to become happier to let a piece of software make decisions on their behalf. This is something that is likely to take time and will only occur as agents show what they are capable of. From a technical perspective, agents need to clearly understand the limit of their responsibility and to act efficiently and safely within these bounds.

ACKNOWLEDGMENTS

The authors would like to thank Xudong Luo of the University of Southampton for his great help on improving earlier versions of this paper. They would also like to thank Yoav Shoham of Stanford University, Robert H. Guttman of Frictionless Commerce, Carles Sierra of the Spanish Artificial Intelligence Research Institute, Chris Preist of Hewlett-Packard Lab, and Dajun Zeng of the University of Arizona, who have provided them with their papers and answered questions about them. Thanks are also given to Nick Gibbins of the University of Southampton for his help on the interaction languages section. Thanks are given to the reviewers for their insightful comments which have significantly improved the quality of the paper.

REFERENCES

- [1] G. Adams, G. Rausser, and L. Simon, "Modelling Multilateral Negotiations: An Application to California Water Policy," *J. Economic Behavior and Organization*, vol. 30, no. 1, pp. 97-111, 1996.
- [2] M.R. Andersson and T.W. Sandholm, "Leveled Commitment Contracting Among Myopic Individually Rational Agents," *Proc. Third Int'l Conf. Multi-Agent Systems*, pp. 26-33, July 1998.
- [3] P. Anthony and N.R. Jennings, "Evolving Bidding Strategies for Multiple Auctions," *Proc. 15th European Conf. Artificial Intelligence*, F. van Harmelen, ed., pp. 178-182, Amsterdam: IOS Press, 2002.
- [4] J. Austin, *How to Do Things with Words*. Oxford, UK, 1962.
- [5] B. Banerjee, A. Biswas, and M. Mundhe, S. Depnath, and S. Sen, "Using Bayesian Networks to Model Agent Relationships," *Applied Artificial Intelligence*, vol. 14, no. 9, pp. 867-879, 2000.
- [6] R. Bapna, P. Goes, and A. Gupta, "Insights and Analyses of Online Auctions," *Comm. ACM*, vol. 44, no. 11, pp. 42-50, 2001.
- [7] P. Barnes-Vieyra and C. Claycomb, "Business-to-Business E-Commerce: Models and Managerial Decisions," *Business Horizons*, pp. 13-20, May-June 2001.
- [8] C. Beam and A. Segev, "Automated Negotiation: A Survey of the State of the Art," CITM Working Paper 96-WP-1022, 1996.
- [9] C. Beam and A. Segev, "Electronic Catalogs and Negotiations," CITM Working Paper 96-WP-1016, 1996.
- [10] T. Berners-Lee, J. Hendler, and O. Lassila, *The Semantic Web*. Scientific American, May 2001.
- [11] M. Bichler, "An Experimental Analysis of Multi-Attribute Auctions," *Decision Support Systems*, vol. 29, pp. 249-268, 2000.
- [12] M. Bichler and M. Kaukal, "Design and Implementation of a Brokerage Service for Electronic Procurement," *Proc. 10th Int'l Workshop Database and Expert Systems Applications*, pp. 618-622, 1999.
- [13] M. Bichler, M. Kaukal, and A. Segev, "Multi-Attribute Auctions for Electronic Procurement," *Proc. First IBM IAC Workshop Internet Based Negotiation Technologies*, 1999.
- [14] M. Bichler and A. Segev, "A Brokerage Framework for Internet Commerce," *Distributed and Parallel Databases*, vol. 7, pp. 133-148, 1999.
- [15] D. Billsus and M. Pazzani, "Learning Collaborative Information Filters," *Proc. 15th Int'l Conf. Machine Learning*, pp. pp. 46-54, 1998.
- [16] A. Birk, "Boosting Cooperation by Evolving Trust," *Applied Artificial Intelligence*, vol. 14, no. 8, pp. 769-784, 2000.
- [17] W. Brenner, R. Zarnekow, H. Witting, *Intelligent Software Agents*. Springer, 1998.
- [18] P.J. Brewer, "Decentralized Computation Procurement and Computational Robustness in a Smart Market," *Economic Theory*, vol. 13, pp. 41-92, 1999.
- [19] A. Byde, C. Preist, and N.R. Jennings, "Decision Procedures for Multiple Auctions," *Proc. First Int'l Joint Conf. Autonomous Agents and Multi-Agent Systems*, pp. 613-620, July 2002.
- [20] C.G. Caplice, "An Optimization Based Bidding Process: A New Framework for Shipper-Carrier Relationships," PhD thesis, School of Engineering, MIT, 1996.
- [21] C. Castelfranchi and R. Falcone, "Trust and Control: A Dialectic Link," *Applied Artificial Intelligence*, vol. 14, no. 8, pp. 799-823, 2000.
- [22] C. Castelfranchi, R. Falcone, B.S. Firozabadi, and Y. Tan, "Guest Editorial," *Applied Artificial Intelligence*, vol. 14, pp. 863-865, 2000.
- [23] B. Chaib-draa, "Industrial Applications of Distributed Artificial Intelligence," *Comm. ACM*, vol. 38, no. 11, pp. 47-53, 1995.
- [24] Y.-K. Che, "Design Competition through Multidimensional Auctions," *RAND J. Economics*, vol. 24, no. 4, pp. 668-680, 1993.
- [25] J. Collins, B. Youngdahl, and S. Jamison,, B. Mobasher, and M. Gini, "A Market Architecture for Multi-Agent Contracting," *Proc. Second Int'l Conf. Autonomous Agents*, K.P. Sycara and M. Wooldridge, eds., pp. 285-292, May 1998.
- [26] V.D. Dang and N.R. Jennings, "Polynomial Algorithms for Clearing Multi-Unit Single Item and Multi-Unit Combinatorial Reverse Auctions," *Proc. 15th European Conf. Artificial Intelligence*, F. Van Harmelen, ed., pp. 23-27, Amsterdam: IOS Press, 2002.
- [27] R. Das, J. Hanson, J. Kephart, and G. Tesauro, "Agent-Human Interactions in the Continuous Double Auction," *Proc. 17th Int'l Joint Conf. Artificial Intelligence*, B. Nebel, ed., vol. 2, pp. 1169-1176, Aug. 2001.
- [28] M. Dastani, N. Jacobs, C.M. Jonker, and J. Treuer, "Modeling User Preferences and Mediating Agents in Electronic Commerce," *Agent Mediated Electronic Commerce*, F. Dignum and C. Sierra, eds., pp. 163-193, 2001.
- [29] M. Devlin and T. Scott, "Using a Speech Technology Agent as an Interface for E-Commerce," *E-Commerce Agents*, J. Liu and Y. Ye, eds., pp. 332-346, 2001.
- [30] W. Dou and D.C. Chou, "A Structural Analysis of Business-to-Business Digital Markets," *Industrial Marketing Management*, vol. 31, no. 2, pp. 165-176, Feb. 2002.

- [31] D. Easley and J. Ledyard, "Theories of Price Formation and Exchange in Double Oral Auctions," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 63-97, 1993.
- [32] A.M. Easwaran and J. Pitt, "An Agent Service Brokering Algorithm for Winner Determination in Combinatorial Auctions," *Proc. 14th European Conf. Artificial Intelligence*, W. Horn, ed., pp. 286-290, 2000.
- [33] P. Faratin, "Automated Service Negotiation Between Autonomous Computational Agents," PhD thesis, Univ. of London, Queen Mary College, 2000.
- [34] P. Faratin, C. Sierra, and N. R. Jennings, "Negotiation Decision Functions for Autonomous Agents," *Int'l J. Robotics and Autonomous System*, vol. 24, nos. 3-4, pp. 159-182, 1998.
- [35] P. Faratin, C. Sierra, and N.R. Jennings, "Using Similarity Criteria to Make Trade-Offs in Automated Negotiations," *Artificial Intelligence*, vol. 142, no. 2, pp. 205-237, 2002.
- [36] T. Finin and Y. Labrou, "KQML as an Agent Communication Language," *Software Agents*, J. Bradshaw, ed., pp. 291-316, Cambridge, Mass.: AAAI Press and MIT Press, 1997.
- [37] R. Fisher and W. Ury, *Getting to Yes: Negotiating an Agreement without Giving In*. Random House, 1981.
- [38] J. Foss, "Brokering the Info-Underworld," *British Telecommunications Eng.*, vol. 17, no. 2, pp. 202-206, 1998.
- [39] *The Grid: Blueprint for a New Computing Infrastructure*. I. Foster and C. Kesselman, eds. Morgan Kaufmann, 1998.
- [40] M.S. Fox, M. Barbuceanu, and R. Teigen, "Agent-Oriented Supply-Chain Management," *The Int'l J. Flexible Manufacturing*, vol. 12, pp. 165-188, 2000.
- [41] D. Friedman, "The Double Auction Market Institution: A Survey," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 3-25, Reading, Mass: Addison-Wesley, 1993.
- [42] Y. Fujishima, K. Leyton-Brown, and Y. Shoham, "Taming the Computational Complexity of Combinatorial Auctions: Optimal and Approximate Approaches," *Proc. 16th Int'l Joint Conf. Artificial Intelligence*, pp. 548-553, 1999.
- [43] I. Gallego, J. Delgado, and J.J. Acebron, "Distributed Models for Brokerage on Electronic Commerce," *Proc. TREC '98*, W. Lamersdorf and M. Merz, eds., pp. 129-140, 1998.
- [44] S. Gamvroulas, D. Polemi, and M. Anagnostou, "A Secure Brokerage Network for Retail Banking Services," *Future Generation Computer Systems*, vol. 16, pp. 423-430, 2000.
- [45] S. Garfinkel, G. Spafford, and D. Russell, *Web Security, Privacy, and Commerce*. O'Reilly, 2001.
- [46] S. Gjerstad and J. Dickhaut, "Price Formation in Double Auction," *Games and Economic Behavior*, vol. 22, pp. 1-29, 1998.
- [47] M. Gleizes et al., "Self-Organization and Learning in Multi-Agent Based Brokerage Services," *Intelligence in Services & Networks*, H. Zuidweg et al. ed., pp. 41-54, 1999.
- [48] R. Glushko, J. Tenenbaum, and B. Meltzer, "An XML Framework for Agent-Based E-Commerce," *Comm. ACM*, vol. 42, no. 3, pp. 106-114, Mar. 1999.
- [49] P. Gmytrasiewicz and E. Durfee, "A Rigorous, Operational Formalization of Recursive Modeling," *Proc. First Int'l Conf. Multi-Agent Systems*, pp. 125-132, 1995.
- [50] D. Gode and S. Sunder, "Lower Bounds for Efficiency of Surplus Extraction in Double Auctions," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 199-219, Addison-Wesley, 1993.
- [51] S.L. Goldman, R.N. Nagel, and K. Preiss, *Agile Competitors and Virtual Organizations: Strategies for Enriching the Customer*. New York: Van Nostrand Reinhold, 1995.
- [52] A. Gomez-Perez and O. Corcho, "Ontology Languages for the Semantic Web," *IEEE Intelligent Systems*, vol. 17, no. 1, pp. 54-60, Jan./Feb. 2002.
- [53] R.H. Guttman and P. Maes, "Agent-Mediated Integrative Negotiation for Retail Electronic Commerce," *Agent Mediated Electronic Commerce*, P. Noriega and C. Sierra, eds., pp. 70-90, 1998.
- [54] R.H. Guttman, A.G. Moukas, and P. Maes, "Agent-Mediated Electronic Commerce: A Survey," *The Knowledge Eng. Rev.*, vol. 13, no. 2, pp. 147-159, 1998.
- [55] S.H. Ha and S.C. Park, "Matching Buyers and Suppliers: An Intelligent Dynamic-Exchange Model," *Intelligent E-Business*, pp. 28-40, July/Aug. 2001.
- [56] *E-Commerce, A World of Opportunity*, R. Hake, ed., Swindon, Wiltshire, UK: The British Computer Soc., 1999.
- [57] F. van Harmelen and I. Horrocks, "FAQs on OIL: The Ontology Inference Layer," *IEEE Intelligent Systems*, vol. 15, no. 6, pp. 69-72, Nov./Dec. 2000.
- [58] M. He, H.F. Leung, and N.R. Jennings, "A Fuzzy Logic Based Bidding Strategy for Autonomous Agents in Continuous Double Auctions," *IEEE Trans. Knowledge and Data Eng.*, to appear.
- [59] J. Hendler and D. McGuinness, "The DARPA Agent Markup Language," *IEEE Intelligent Systems*, vol. 15, no. 6, pp. 72-73, Nov./Dec. 2000.
- [60] M. Huhns, L. Stephens, and N. Ivezic, "Automating Supply-Chain Management," *Proc. First Int'l Joint Conf. Autonomous Agents and Multi-Agent Systems*, pp. 1017-1024, July 2002.
- [61] M.N. Huhns and M.P. Singh, "Ontologies for Agents," *IEEE Internet Computing*, vol. 1, no. 6, pp. 81-83, 1997.
- [62] T. Ito, N. Fukuta, T. Shintani, and K. Sycara, "Biddingbot: A Multiagent Support System for Cooperative Bidding in Multiple Auctions," *Proc. Fourth Int'l Conf. Multi-Agent Systems*, pp. 399-400, 2000.
- [63] N.R. Jennings, "On Agent-Based Software Engineering," *Artificial Intelligence*, vol. 117, no. 2, pp. 277-296, 2000.
- [64] N.R. Jennings, "An Agent-Based Approach for Building Complex Software Systems," *Comm. ACM*, vol. 44, no. 4, pp. 35-41, 2001.
- [65] N.R. Jennings, P. Faratin, A.R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge, "Automated Negotiation: Prospects, Methods and Challenges," *Int. J. Group Decision and Negotiation*, vol. 10, no. 2, pp. 199-215, 2001.
- [66] N.R. Jennings, T.J. Norman, P. Faratin, P. O'Brien, and B. Odgers, "Autonomous Agents for Business Process Management," *Applied Artificial Intelligence*, vol. 14, no. 2, pp. 145-189, 2000.
- [67] N.R. Jennings, K. Sycara, and M. Wooldridge, "A Roadmap of Agent Research and Development," *Autonomous Agent and Multi-Agent Systems*, vol. 1, pp. 7-38, 1998.
- [68] N.R. Jennings, K. Woghiren, and S. Osborn, "Interacting Agents—The Way Forward for Agent-Mediated Electronic Commerce," technical report, 2000. Lost Wax white paper, available at <http://www.lostwax.com/lostwax1/opinions/opinions1.html>.
- [69] N.R. Jennings and M. Wooldridge, "Applying Agent Technology," N.R. Jennings and M. Wooldridge, eds., *Agent Technology: Foundations, Applications, and Markets*, Berlin, Germany: Springer-Verlag, 1998.
- [70] J. Jung and G. Jo, "Brokerage between Buyer and Seller Agents Using Constraint Satisfaction Problem Models," *Decision Support Systems*, vol. 28, pp. 293-304, 2000.
- [71] *Frontiers of Electronic Commerce*, chapter 16, R. Kalakota and A.B. Whinston, eds., pp. 595-628, Addison-Wesley, 1996.
- [72] R. Karp, "Reducibility among Combinatorial Problems," *Complexity of Computer Computations*, R. Miller and J. Thatcher, eds., pp. 85-103, Cambridge, Mass.: MIT Press, 1972.
- [73] P. Keskinocak and R. Goodwin et al., "Decision Support for Managing an Electronic Supply Chain," *Electronic Commerce Research*, vol. 1, nos. 1-2, pp. 15-31, 2001.
- [74] R. Khosla and S. Kitjonthawonkul, "A Human-Centered Agent-Based Architecture for Electronic Brokerage," *Soft Computing*, vol. 6, pp. 405-411, 2001.
- [75] M.T. Kone, A. Shimazu, and T. Nakajima, "The State of the Art in Agent Communication Languages," *Knowledge and Information Systems*, vol. 2, pp. 259-284, 2000.
- [76] R. Kowalczyk, "On Negotiation as a Distributed Fuzzy Constraint Satisfaction Problem," *Proc. DEXA E-Negotiation Workshop*, pp. 631-637, 2000.
- [77] R. Kowalczyk and V. Bui, "On Fuzzy E-Negotiation Agents: Autonomous Negotiation with Incomplete and Imprecise Information," *Proc. DEXA e-Negotiation Workshop*, 2000.
- [78] S. Kraus, *Strategic Negotiation in Multi-Agent Environments*. Cambridge, Mass.: MIT Press, 2001.
- [79] D. Kreps, *Game Theory and Economic Modeling*. Oxford University Press, 1990.
- [80] B. Krulwich, "The Bargainfinder Agent: Comparison Price Shopping on the Internet," *Bots, and Other Internet Beasts*, J. Williams, ed., pp. 257-263, Macmillan Computer Publishing, 1996.
- [81] K. Kumar, "Technology for Supporting Supply-Chain Management," *Comm. ACM*, vol. 44, no. 6, pp. 58-61, 2001.
- [82] Y. Labrou, T. Finin, and Y. Peng, "Agent Communication Languages: The Current Landscape," *IEEE Intelligent Systems*, vol. 14, no. 2, pp. 45-52, Mar./Apr. 1999.
- [83] S.S. Laurent, *XML: Element of Style*. McGraw-Hill, 2000.

- [84] H.L. Lee and C. Billington, "The Evolution of Supply-Chain-Management Models and Practice at Hewlett-Packard," *Interface*, vol. 25, no. 5, pp. 42-63, Sep-Oct. 1995.
- [85] J.-J. Lee and R. McCartney, "Predicting User Actions Using Interface Agents with Individual User Models," *Approaches to Intelligent Agents*, H. Nakashima and C. Zhang, eds., pp. 154-169, 1999.
- [86] T. Liang and J. Huang, "A Framework for Applying Intelligent Agents to Support Electronic Trading," *Decision Support Systems*, vol. 28, pp. 305-317, 2000.
- [87] J. Liu, K. Wong, and K. Hui, "Discovering User Behavior Patterns in Personalized Interface Agents," *Intelligent Data Eng. and Automated Learning*, K.S. Leung, L.-W. Chan, and H. Meng, eds., pp. 398-403, 2000.
- [88] J. Liu and Y. Ye, "Introduction to E-Commerce Agents: Marketplace Solutions, Security Issues, and Supply and Demand," *E-Commerce Agents: Marketplace Solutions, Security Issues, and Supply and Demand*, J. Liu and Y. Ye, eds., pp. 1-6, 2001.
- [89] X. Luo, N.R. Jennings, N. Shadbolt, H.F. Leung, and J.H.M. Lee, "A Fuzzy Constraint Based Model for Bilateral, Multi-Issue Negotiation in Semi-Competitive Environments," *Artificial Intelligence*, vol. 148, nos. 1-2, 2003.
- [90] M. Ma, "Agents in E-Commerce," *Comm. ACM*, vol. 42, no. 3, pp. 79-80, Mar. 1999.
- [91] P. Maes, "Agents That Reduce Work and Information Overload," *Comm. ACM*, vol. 37, no. 7, pp. 31-40, 1994.
- [92] M.T. Martinez, P. Foulletier, K.H. Park, and J. Favrel, "Virtual Enterprise—Organisation, Evolution and Control," *Int'l J. Production Economics*, vol. 74, pp. 225-238, 2001.
- [93] P. Massotte, "New Concepts and Approaches for the Management of Complex Production Systems," *Proc. Int'l Conf. Industrial Eng.*, 1993.
- [94] N. Matos, C. Sierra, and N.R. Jennings, "Determining Successful Negotiation Strategies: An Evolutionary Approach," *Proc. Third Int'l Conf. Multi-Agent Systems*, pp. 182-189, 1998.
- [95] R.P. McAfee and J. McMillan, "Auctions and Bidding," *J. Economic Literature*, vol. 25, pp. 699-738, June 1987.
- [96] P. McDermott, "Building Trust into Online Business," *Network Security*, vol. 10, pp. 10-12, 2000.
- [97] L. Meade, J. Sarkis, and D. Liles, "Justifying Strategic Alliances: A Prerequisite for Virtual Enterprising," *OMEGA: The Int'l J. Management Science*, vol. 25, no. 1, pp. 29-42, 1996.
- [98] P. Milgrom, "Auctions and Bidding: A Primer," *J. Economic Perspectives*, pp. 3-22, 1989.
- [99] P. Milgrom and J.W. Robert, "A Theory of Auctions and Competitive Bidding," *Econometrica*, pp. 1089-1122, 1982.
- [100] L. Mui, A. Halberstadt, and M. Mohtashemi, "Notions of Reputation in Multi-Agents Systems: A Review," *Proc. First Int'l Joint Conf. Autonomous Agents and Multi-Agent Systems*, pp. 280-287, July 2002.
- [101] J.P. Müller, "Architectures and Applications of Intelligent Agents: A Survey," *The Knowledge Eng. Rev.*, vol. 13, no. 4, pp. 353-380, 1998.
- [102] *Intelligent Software Agents*, R. Murch and T. Johnson, eds. chapter 14, pp. 131-137, Upper Saddle River, N.J.: Prentice Hall PTR, 1999.
- [103] H.S. Nwana and D.T. Ndumu, "A Perspective on Software Agents Research," *The Knowledge Eng. Rev.*, vol. 14, no. 2, pp. 125-142, 1999.
- [104] P. O'Brien and R. Nicol, "FIPA—Towards a Standard for Software Agents," *BT Technology J.*, vol. 16, no. 3, pp. 51-59, July 1998.
- [105] D.E. O'Leary, D. Kuokka, and R. Plant, "Artificial Intelligence and Virtual Organizations," *Comm. ACM*, vol. 40, no. 1, pp. 52-59, Jan. 1997.
- [106] S. Park, E. Durfee, and W. Birmingham, "Advantages of Strategic Thinking in Multiagent Contracts (a Mechanism and Analysis)," *Proc. Second Int'l Conf. Multi-Agent Systems*, pp. 259-266, 1996.
- [107] S. Park, E. Durfee, and W. Birmingham, "An Adaptive Agent Bidding Strategy Based on Stochastic Modeling," *Proc. Third Int'l Conf. Autonomous Agents*, pp. 147-153, 1999.
- [108] D.C. Parkes, L.H. Ungar, and D.P. Foster, "Accounting for Cognitive Costs in On-Line Auction Design," *Agent Mediated Electronic Commerce*, P. Noriega and C. Sierra, eds., pp. 25-40, 1998.
- [109] H.V. Parunak, "Industrial and Practical Applications of DAL," *Multi-Agent Systems*, G. Weib, ed., Cambridge, Mass.: MIT Press, 1998.
- [110] C. Preist, C. Bartolini, and I. Phillips, "Algorithm Design for Agents Which Participate in Multiple Simultaneous Auctions," *Agent Mediated Electronic Commerce III*, F. Dignum and U. Cortes, eds., pp. 139-154, Springer-Verlag, 2001.
- [111] C. Preist, A. Bye, and C. Bartolini, "Economic Dynamics of Agents in Multiple Auctions," *Proc. Fifth Int'l Conf. Autonomous Agents*, pp. 545-551, 2001.
- [112] R.J. Rabelo, L.M. Camarinha-Matos, and H. Afsarmanesh, "Multi-Agent-Based Agile Scheduling," *Robotics and Autonomous Systems*, vol. 27, pp. 15-28, 1999.
- [113] S. Rassenti, V. Smith, and R. Bulfin, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell J. Economics*, vol. 13, no. 2, pp. 402-417, 1982.
- [114] R.A. Ribeiro, "Fuzzy Multiple Attribute Decision Making: A Review and New Preference Elicitation Techniques," *Fuzzy Sets and Systems*, vol. 78, pp. 155-181, 1996.
- [115] J.S. Rosenschein and G. Zlotkin, *Rules of Encounter*. The MIT Press, 1994.
- [116] D. de Roure, N.R. Jennings, and N. Shadbolt, "The Semantic Grid: A Future E-Science Infrastructure," *Int'l J. Concurrency and Computation: Practice and Experience*, vol. 15, no. 11, 2003.
- [117] J. Rust, J.H. Miller, and R. Palmer, "Behavior of Trading Automata in a Computerized Double Auction Market," *The Double Auction Market: Institutions, Theories, and Evidence*, D. Friedman and J. Rust, eds., pp. 155-198, Addison-Wesley, 1991.
- [118] A. Sadrieh, *The Alternating Double Auction Market: A Game Theoretic and Experimental Investigation*. Springer, 1998.
- [119] T. Sandholm, "An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations," *Proc. First Int'l Conf. Artificial Intelligence*, pp. 256-262, 1993.
- [120] T. Sandholm, "Unenforced E-Commerce Transactions," *IEEE Internet Computing*, vol. 1, no. 6, pp. 47-54, 1997.
- [121] T. Sandholm, "Automated Negotiation," *Comm. ACM*, vol. 42, no. 3, pp. 84-85, 1999.
- [122] T. Sandholm, "Distributed Rational Decision Making," *Multiagent Systems*, G. Weiss, ed., pp. 201-258, The MIT Press, 1999.
- [123] T. Sandholm, "eMediator: A Next Generation Electronic Commerce Server," *Proc. AAAI '99 Workshop*, 1999.
- [124] T. Sandholm, "Algorithm for Optimal Winner Determination in Combinatorial Auctions," *Artificial Intelligence*, vol. 135, pp. 1-54, 2002.
- [125] J. Sarkis and R.P. Sundarraj, "Evolution of Brokering: Paradigms in E-Commerce Enabled Manufacturing," *Int. J. Production Economics*, vol. 75, pp. 21-31, 2002.
- [126] C.M. Sashi and B. O'Leary, "The Role of Internet Auctions in the Expansion of B2B Markets," *Industrial Marketing Management*, vol. 31, pp. 103-110, 2002.
- [127] M. Schillo, P. Funk, and M. Rovatsos, "Using Trust for Detecting Deceitful Agents in Artificial Societies," *Applied Artificial Intelligence*, vol. 14, no. 8, pp. 825-848, 2000.
- [128] B. Schmid and M. Lindemann, "Elements of a Reference Model for Electronic Markets," *Proc. 31st Hawaii Int'l Conf. System Sciences*, 1998.
- [129] J. Searle, *Speech Acts: An Essay in the Philosophy of Language*. Cambridge Univ. Press, 1970.
- [130] S. Sen, A. Biswas, and S. Debnath, "Believing Others: Pros and Cons," *Proc. Fourth Int'l Conf. Multi-Agent Systems*, pp. 279-285, 2000.
- [131] U. Shardanand and P. Maes, "Social Information Filtering: Algorithms for Automating Word of Mouth," *Proc. Computer Factors in Computing Systems*, pp. 210-217, 1995.
- [132] A. Sharma, "Trends in Internet-Based Business-to-Business Marketing," *Industrial Marketing Management*, vol. 31, pp. 77-84, 2002.
- [133] M.J. Shaw, "Electronic Commerce: State of Art," *Handbook on Electronic Commerce*, M. Shaw, R. Blanning, T. Strader, and A. Whinston, eds., chapter 1, pp. 3-24, Springer, 2000.
- [134] O. Shehory and S. Kraus, "Methods for Task Allocation via Agent Coalition Formation," *Artificial Intelligence*, vol. 101, pp. 165-200, 1998.
- [135] C. Sierra and F. Dignum, "Agent-Mediated Electronic Commerce: Scientific and Technological Roadmap," *Agent Mediated Electronic Commerce*, F. Dignum and C. Sierra, eds., pp. 1-18, Springer, 2001.
- [136] C. Sierra, P. Faratin, and N.R. Jennings, "A Service-Oriented Negotiation Model between Autonomous Agents," *Lecture Notes in Artificial Intelligence*, vol. 1237, pp. 17-35, 1997.

- [137] C. Sierra, N.R. Jennings, P. Noriega, and S. Parsons, "A Framework for Argumentation-Based Negotiation," *Proc. Fourth Int'l Workshop on Agent Theories, Architectures, and Languages*, pp. 177-192, 1997.
- [138] S.N. Singh, "Web Home Pages as Advertisements," *Comm. ACM*, vol. 42, no. 8, pp. 91-98, 1999.
- [139] R. Smith, "The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver," *IEEE Trans. Computers*, vol. 29, no. 12, pp. 1104-1113, Dec. 1980.
- [140] S.S. Standifird, "Reputation and E-Commerce: eBay Auctions and the Asymmetrical Impact of Positive and Negative Ratings," *J. Management*, vol. 27, pp. 279-295, 2001.
- [141] M. Subramani and E. Walden, "Economic Returns to Firms from Business-to-Business Electronic Commerce Initiatives: An Empirical Examination," *Proc. 21st Int'l Conf. Information Systems*, pp. 229-241, 2000.
- [142] R. Sun, B. Chu, R. Wilhelm, and J. Yao, "A CSP-Based Model for Integrated Supply Chain," *Proc. AAAI-99 Workshop AI for Electronic Commerce*, 1999.
- [143] G.W. Tan, M.J. Shaw, and W. Fulkerson, "Web-Based Global Supply Chain Management," *Handbook on Electronic Commerce*, M. Shaw, R. Blanning, T. Strader, and A. Whinston, eds., chapter 22, pp. 457-478, Springer, 2000.
- [144] J. Teich, H. Wallenius, and J. Wallenius, "Multiple-Issue Auction and Market Algorithms for the World Wide Web," *Decision Support Systems*, vol. 26, pp. 49-66, 1999.
- [145] L.G. Telsler, "A Theory of Self-Enforcing Agreements," *J. Business*, vol. 53, no. 1, pp. 27-44, 1980.
- [146] G. Tesauo and J. Bredin, "Strategic Sequential Bidding in Auctions Using Dynamic Programming," *Proc. First Int'l Joint Conf. Autonomous Agents and Multi-Agent Systems*, pp. 591-598, July 2002.
- [147] R. Till, "Transforming the Way We Do Business," *Electronic Commerce*, T. Nash, ed., pp. 9-12, 1998.
- [148] J. Tirole, *The Theory of Industrial Organisation*. Cambridge, Mass.: The MIT Press, 1998.
- [149] M. Tsvetovat and K. Sycara, "Customer Coalitions in the Electronic Marketplace," *Proc. Fourth Int'l Conf. Autonomous Agents*, pp. 263-264, 2000.
- [150] S. Tully, "The B2B Tool that Really is Changing the World," *Fortune*, pp. 135-145, Mar. 2000.
- [151] A. Tuma, "Configuration and Coordination of Virtual Production Networks," *Int'l J. Production Economics*, vols. 56-57, pp. 641-648, 1998.
- [152] *Electronic Commerce: A Managerial Perspective*, E. Turban, J. Lee, D. King, and H.M. Chung, eds. Prentice Hall, 1999.
- [153] W. Vickrey, "Counterspeculation, Auctions and Competitive Sealed Tenders," *J. Finance*, vol. 16, pp. 8-37, 1961.
- [154] J. Vidal and E. Durfee, "The Impact of Nested Agent Models in an Information Economy," *Proc. Second Int'l Conf. Multi-Agent Systems*, pp. 377-384, 1996.
- [155] S. Vries and R.V. Vohra, "Combinatorial Auctions: A Survey," technical report, 2001. <http://www-m9.mathematik.tu-muenchen.de/dm/homepages/devries/>.
- [156] N. Vulkan and N.R. Jennings, "Efficient Mechanisms for the Supply of Services in Multi-Agent Environments," *Int'l J. Decision Support Systems*, vol. 28, nos. 1-2, pp. 5-19, 2000.
- [157] W. Walsh and M. Wellman, "Modeling Supply Chains Formation in Multiagent Systems," *Proc. IJCAI-99 Workshop Agent-Mediated Electronic Commerce*, 1999.
- [158] G. Weiss, "Agent Orientation in Software Engineering," *The Knowledge Eng. Rev.*, vol. 16, no. 4, pp. 349-373, 2001.
- [159] W. Wen and F. Mizoguchi, "An Authorization-Based Trust Model for Multiagent Systems," *Applied Artificial Intelligence*, vol. 14, no. 9, pp. 909-925, 2000.
- [160] E. Wolfstetter, *Topics in Microeconomics: Industrial Organization, Auctions, and Incentives*, chapter 8, pp. 182-242, Cambridge Univ. Press, 1999.
- [161] H.C. Wong and K. Sycara, "Adding Security and Trust to Multiagent Systems," *Applied Artificial Intelligence*, vol. 14, no. 9, pp. 927-941, 2000.
- [162] M. Wooldridge and N.R. Jennings, "Intelligent Agents: Theory and Practice," *The Knowledge Eng. Rev.*, vol. 10, no. 2, pp. 115-152, 1995.
- [163] P.R. Wurman, "Dynamic Pricing in the Virtual Marketplace," *IEEE Internet Computing*, vol. 5, pp. 36-42, Mar./Apr. 2001.

- [164] P.R. Wurman, M. Wellman, and W. Walsh, "The Michigan Internet Auctionbot: A Configurable Auction Server for Human and Software Agents," *Proc. Second Int'l Conf. Autonomous Agents*, pp. 301-308, 1998.
- [165] J. Yamamoto and K. Sycara, "A Stable and Efficient Buyer Coalition Formation Scheme for E-Marketplaces," *Proc. Fifth Int'l Conf. Autonomous Agents*, pp. 576-583, 2001.
- [166] Y. Ye, J. Liu, and A. Moukas, "Agents in Electronic Commerce," *Electronic Commerce Research*, vol. 1, nos. 1-2, pp. 9-14, Feb./Mar. 2001.
- [167] E. Young, "Web Marketplaces That Really Work," technical report, Fortune/CNET Tech., 2001.
- [168] G. Zacharia and P. Maes, "Trust Management through Reputation Mechanisms," *Applied Artificial Intelligence*, vol. 14, no. 9, pp. 881-907, 2000.
- [169] G. Zacharia, A. Moukas, and P. Maes, "Collaborative Reputation Mechanisms for Electronic Marketplaces," *Decision Support Systems*, vol. 29, pp. 371-388, 2000.
- [170] D. Zeng, "Managing Flexibility for Inter-Organizational Electronic Commerce," *Electronic Commerce Research*, vol. 1, nos. 1-2, pp. 33-51, 2001.
- [171] D. Zeng and K. Sycara, "Bayesian Learning in Negotiation," *Int'l J. Human-Computer Studies*, vol. 48, pp. 125-141, 1998.



Minghua He received the master's degree in computer science from Southwest China Normal University in China in 1999. She is a PhD candidate in the Department of Electronics and Computer Science at the University of Southampton, United Kingdom. Her research interests center around the study of artificial intelligence, especially the area of agent-based automated negotiation for E-commerce.



Nicholas R. Jennings is a professor of computer science in the Department of Electronics and Computer Science at Southampton University where he carries out basic and applied research in agent-based computing. He is head of the Intelligence, Agents, Multimedia Group and is also the chief scientific officer for Lost Wax. Professor Jennings helped pioneer the use of agent-based techniques for real-world applications and he has also made theoretical contributions in the areas of automated negotiation and auctions, cooperative problem solving, and socially rational decision making. He has published some 150 articles on various facets of agent-based computing, holds two patents (three more pending), written one monograph, and coedited five books. He is in the top 250 most cited computer scientists (according to the citeseer digital library). He was the recipient of the Computers and Thought Award (the premier award for a young AI scientist) in 1999 for his contributions to practical agent architectures and applications of multiagent systems (this is the first time in the Award's 30 year history that it has been given to someone based in Europe), and the recipient of an IEE Achievement Medal in 2000 for his work on agent-based computing.



Ho-Fung Leung received the BSc (Hon.) and M.Phil degrees in computer science from The Chinese University of Hong Kong. He received the PhD degree in 1992 from the Imperial College of Science, Technology, and Medicine. He is an associate professor in the Department of Computer Science and Engineering at The Chinese University of Hong Kong. Dr. Leung joined the Department of Computer Science at The Chinese University of Hong Kong in 1992. His current research interests include constraint satisfaction, multiagent systems, and logic programming. He is a participating member of the Centre for the Advancement of E-Commerce Technologies (AECT) of the Engineering Faculty.