

Agent-mediated Electronic Commerce: A Survey

Robert H. Guttman, Alexandros G. Moukas, and Pattie Maes

Software Agents Group
MIT Media Laboratory
20 Ames Street, E15-305
Cambridge, MA 02139

{guttman,moux,pattie}@media.mit.edu
<http://ecommerce.media.mit.edu>

Abstract

Software agents help automate a variety of tasks including those involved in buying and selling products over the Internet. This paper surveys several of these agent-mediated electronic commerce systems by describing their roles in the context of a Consumer Buying Behavior (CBB) model. The CBB model we present augments traditional marketing models with concepts from Software Agents research to accommodate electronic markets. We then discuss the variety of Artificial Intelligence techniques that support agent mediation and conclude with future directions of agent-mediated electronic commerce research.

1. Introduction

Software agents are programs to which one can delegate (aspects of) a task. They differ from “traditional” software in that they are personalized, continuously running and semi-autonomous. These qualities make agents useful for a wide variety of information and process management tasks [1]. It should come as no surprise that these same qualities are particularly useful for the information-rich and process-rich environment of electronic commerce.

Electronic commerce encompasses a broad range of issues including security, trust, reputation, law, payment mechanisms, advertising, ontologies, on-line catalogs, intermediaries, multimedia shopping experiences, and back-office management. Agent technologies can be applied to any of these areas where a personalized, continuously running, semi-autonomous behavior is desirable. However, certain characteristics will determine to what extent agent technologies are appropriate.

For example, how much time or money could be saved if a certain process was partially automated (e.g., comparing products from multiple merchants)? How easy is it to express your preferences for the task (e.g., shopping for a gift)? What are the risks of an agent making a sub-optimal transaction decision (e.g., making stock market buying and selling decisions or buying a car)? What are the consequences for missed

opportunities (e.g., not being able to effectively monitor new job postings)? Generally, the more time and money that can be saved through automation, the easier it is to express preferences, the lesser the risks of making sub-optimal transaction decisions, and the greater the loss for missed opportunities, the more appropriate it is to employ agent technologies in electronic commerce.

Software agents will play an increasing variety of roles as mediators in electronic commerce [2]. This paper explores these roles, their supporting technologies, and how they relate to electronic commerce in its three main forms: business-to-business, business-to-consumer, and consumer-to-consumer transactions (with an emphasis on the latter two).

2. Roles of Agents as Mediators in Electronic Commerce

It is useful to explore the roles of agents as mediators in electronic commerce in the context of a common model. The model we present stems from traditional marketing Consumer Buying Behavior (CBB) research and comprises the actions and decisions involved in buying and using goods and services. However, we augment traditional CBB models with concepts from Software Agents research to accommodate electronic markets.

Although CBB research covers many areas, it is important to recognize its limitations upfront. For example, CBB research focuses primarily on retail markets (although many CBB concepts pertain to business-to-business and consumer-to-consumer markets as well) [3, 4]. Even within retail, not all shopping behaviors are captured (e.g., impulse purchasing). Also, as mentioned earlier, electronic commerce covers a broad range of issues, some of which are beyond the scope of a CBB model (e.g., back-office management and other merchant issues). Nevertheless, the CBB model is a powerful tool to help us understand the roles of agents as mediators in electronic commerce.

2.1. Consumer Buying Behavior Model

There are several descriptive theories and models that attempt to capture consumer buying behavior – e.g., the Nicosia model [5], the Howard-Sheth model [6], the Engel-Blackwell model [7], the Bettman information-processing model [8], and the Andreasen model [9]. Although different, these models all share a similar list of six fundamental stages guiding consumer buying behavior. These six stages also elucidate where agent technologies apply to the consumer shopping experience and allow us to more formally categorize existing agent-mediated electronic commerce systems [10]:

1. Need Identification

This stage characterizes the consumer becoming aware of some unmet need. Within this stage, the consumer can be stimulated through product information. This stage is called *Problem Recognition* in the Engel-Blackwell model [7].

2. Product Brokering

This stage comprises the retrieval of information to help determine *what* to buy. This encompasses the evaluation of product alternatives based on consumer-provided criteria. The result of this stage is called the “consideration set” of products.

3. Merchant Brokering

This stage combines the “consideration set” from the previous stage with merchant-specific information to help determine *who* to buy from. This includes the evaluation of merchant alternatives based on consumer-selected criteria (e.g., price, warranty, availability, delivery time, reputation, etc.). The Nicosia model merges both brokering stages into one *Search Evaluation* stage [5]. The Engel-Blackwell model dissects these two stages orthogonally into *Information Search* and *Evaluation of Alternatives* stages [7].

4. Negotiation

This stage is about *how* to determine the terms of the transaction. Negotiation varies in duration and complexity depending on the market. In traditional retail markets, prices and other aspects of the transaction are often fixed leaving no room for negotiation. In other markets (e.g., stocks, automobile, fine art, local markets, etc.), the negotiation of price or other aspects of the deal are integral to product and merchant brokering. Traditional CBB models do not identify this stage explicitly, but the conclusion of the Negotiation stage is comparable to the *Choice* or *Decision* stage found in other models [5, 7].

5. Purchase and Delivery

The purchase and delivery of a product can either

signal the termination of the negotiation stage or occur sometime afterwards (in either order). In some cases, the available payment options (e.g., cash only) or delivery options may influence product and merchant brokering.

6. Service and Evaluation

This post-purchase stage involves product service, customer service, and an evaluation of the satisfaction of the overall buying experience and decision. The nature of this stage (and others) depends upon for whom the product was purchased.

As with most models, these stages represent an approximation and simplification of complex behaviors. As noted, CBB stages often overlap and migration from one to another can be non-linear and iterative.

From this CBB perspective, we can identify the roles for agents as mediators in electronic commerce. The personalized, continuously-running, semi-autonomous nature of agents make them well-suited for mediating those consumer behaviors involving information filtering and retrieval, personalized evaluations, complex coordinations, and time-based interactions. Specifically, these roles correspond (most notably) to the Product Brokering, Merchant Brokering, and Negotiation stages of the Consumer Buying Behavior model.

Table 1 lists the six CBB stages and shows where several representative agent systems fall within this space. The rest of this section expounds the three agent-centric stages of the CBB model with examples.

2.2. Product Brokering

The *Product Brokering* stage of the CBB model is where consumers determine *what* to buy. This occurs after a need has been identified (i.e., in the Need Identification stage) and is achieved through a critical evaluation of retrieved product information. Table 1 shows several agent systems that lower consumers’ search costs [11] when deciding which products best meet his or her personal criteria: PersonaLogic, Firefly, and Tete-a-Tete.

PersonaLogic [12] is a tool that enables consumers to narrow down the products that best meet their needs by guiding them through a large product feature space. The system filters out unwanted products within a given domain by allowing shoppers to specify constraints on a product’s features. A constraint satisfaction engine then returns an ordered list of only those products that satisfy all of the hard constraints.

Like PersonaLogic, **Firefly** services [13, 14] help consumers find products. However, instead of filtering

	Persona Logic	Firefly	Bargain Finder	Jango	Kasbah	Auction Bot	Tete-a-Tete
1. Need Identification	Only a few primitive event-alerting tools (e.g., Amazon.com's "Eyes" program) help anticipate consumers' needs and provide paths into the subsequent CBB stages. However, systems like Firefly can alert a consumer with product recommendations when consumers with similar interests purchase specific products.						
2. Product Brokering	X	X		X			X
3. Merchant Brokering			X	X	X		X
4. Negotiation					X	X	X
5. Purchase and Delivery	Post-purchase evaluation usually includes feedback about two distinct elements of the shopping process: product brokering and merchant brokering. Traditionally, customer remarks are accessible (and used) by either the marketing staff of manufacturers or the customer satisfaction staff of merchants. However, agent-based distributed trust and reputation mechanisms (e.g., Kasbah's Better Business Bureau) enable customers to share and combine their experiences and use merchant and product reputations as additional aspects of brokering and negotiation.						
6. Product Service & Eval.							

Table 1 : Roles and Examples of Agent Systems as Mediators in Electronic Commerce

products based on features, Firefly recommends products via a "word of mouth" recommendation mechanism called automated collaborative filtering (ACF). ACF first compares a shopper's product ratings with those of other shoppers. After identifying the shopper's "nearest neighbors" (i.e., users with similar tastes), ACF recommends products that they rated highly but which the shopper has not yet rated, potentially resulting in serendipitous finds. Essentially, Firefly uses the opinions of like-minded people to offer recommendations. The system is currently used to recommend commodity products such as music and books.

2.3. Merchant Brokering

Whereas the Product Brokering stage compares product alternatives, the *Merchant Brokering* stage compares merchant alternatives.

Andersen Consulting's **BargainFinder** was the first shopping agent for on-line price comparisons [15]. Given a specific product, BargainFinder requests its price from each of nine different merchant Web sites using the same request as from a Web browser. Although a limited proof-of-concept system, BargainFinder offers valuable insights into the issues involved in price comparisons in the on-line world. For example, one third of the on-line CD merchants accessed by BargainFinder blocked all of its price requests. One reason for this was that merchants inherently do not want to compete on price alone. Value-added services that merchants offered on their Web sites were being bypassed by BargainFinder and therefore not considered in the consumer's buying decision. However, it was also the case that Andersen Consulting received requests from an equal number of little-known mer-

chants who wanted to be included in BargainFinder's price comparison. In short, companies competing on price and/or desiring more exposure wanted to be included, the others didn't.

Jango [16, 17] can be viewed as an advanced BargainFinder. The original Jango version "solved" the merchant blocking issue by having the product requests originate from each consumer's Web browser instead of from a central site as in BargainFinder. This way, requests to merchants from a Jango-augmented Web browser appeared as requests from "real" customers. This kind of "aggressive interoperability" makes it convenient for consumers to shop for commodity products but does not leave merchants with many options. If merchants provide public on-line catalogs, they can be accessed by agents whether merchants want this or not.

Jango's modus operandi is simple: once a shopper has identified a specific product, Jango can simultaneously query merchant sites (from a list now maintained by Excite, Inc.) for its price. These results allow a consumer to compare merchant offerings on price.

The MIT Media Lab's **Kasbah** [18, 19] is an on-line, multiagent classified ad system. A user wanting to buy or sell a good creates an agent, gives it some strategic direction, and sends it off into a centralized agent marketplace. Kasbah agents proactively seek out potential buyers or sellers and negotiate with them on behalf of their owners. Each agent's goal is to complete an acceptable deal, subject to a set of user-specified constraints such as a desired price, a highest (or lowest) acceptable price, and a date by which to complete the transaction. The latest version of Kasbah incorporates a distributed trust and reputation mechanism called the Better Business Bureau. Upon the completion of a

transaction, both parties may rate how well the other party managed their half of the deal (e.g., accuracy of product condition, completion of the transaction, etc.). Agents can then use these ratings to determine if they should negotiate with agents whose owners fall below a user-specified reputation threshold.

2.4. Negotiation

From our CBB perspective, the *Negotiation* stage is where the price or other terms of the transaction are determined.¹ Examples of where we see negotiation used in commerce include stock markets (e.g., NYSE and NASDAQ), fine art auction houses (e.g., Sotheby's and Christie's), flower auctions (e.g., Aalsmeer, Holland), and various ad-hoc haggling (e.g., automobile dealerships and commission-based electronics stores).

The benefit of dynamically negotiating a price for a product instead of fixing it is that it relieves the merchant from needing to determine the value of the good a priori. Rather, this burden is pushed into the marketplace itself. A result of this is that limited resources are allocated fairly – i.e., to those who value them most.

However, there are impediments to using negotiation. In the physical world, certain types of auctions require that all parties be geographically co-located, for example, in auction houses. Also, negotiating may be too complicated or frustrating for the average consumer. For instance, this sentiment inspired Saturn automobile dealerships to switch from price negotiation to fixed-price in order to appease its customers. Finally, some negotiation protocols occur over an extended period of time which does not cater to impatient or time-constrained consumers. In general, real-world negotiations accrue transaction costs that may be too high for either consumers or merchants [20].

Fortunately, many of these impediments disappear in the digital world. For example, OnSale [21] and eBay's AuctionWeb [22] are two popular Web sites that sell refurbished and second-hand products using a choice of auction protocols. Unlike auction houses, these sites do not require that participants be geographically co-located. However, these sites still require that consumers manage their own negotiation strategies over an extended period of time. This is where agent technologies come in.

Table 1 shows several representative agent systems

1. Like the term "agent", there is no consensus on the definition of the term "negotiation." Economists, game theorists, business managers, political scientists, and artificial intelligence researchers each provide unique perspectives on its meaning. We offer a broad definition of "negotiation" in section 3.3.

that assist customers in negotiating the terms of a transaction: AuctionBot, Kasbah, and Tete-a-Tete.

AuctionBot [23, 24] is a general purpose Internet auction server at the University of Michigan. AuctionBot users create new auctions to sell products by choosing from a selection of auction types and then specifying its parameters (e.g., clearing times, method for resolving bidding ties, the number of sellers permitted, etc.). Buyers and sellers can then bid according to the multi-lateral distributive negotiation protocols of the created auction. In a typical scenario, a seller would bid a reservation price after creating the auction and let AuctionBot manage and enforce buyer bidding according to the auction protocol and parameters.

What makes AuctionBot different from most other auction sites, however, is that it provides an application programmable interface (API) for users to create their own software agents to autonomously compete in the AuctionBot marketplace. Such an API provides a semantically sound interface to the marketplace unlike the "wrapper" technologies discussed in sections 3.1 and 3.4. However, as with the Fishmarket Project [25, 26], it is left to the users to encode their own bidding strategies. Fishmarket is not currently being used as a real-world system, but it has hosted tournaments to compare opponents' hand-crafted bidding strategies [27] along the lines of Axelrod's prisoner's dilemma tournaments [28].

Kasbah, as described earlier, is a Web-based multi-agent classified ad system where users create buying and selling agents to help transact products. These agents automate much of the Merchant Brokering and Negotiation CBB stages for both buyers and sellers.

Negotiation in Kasbah is straightforward. After buying agents and selling agents are matched, the only valid action in the negotiation protocol is for buying agents to offer a bid to selling agents with no restrictions on time or price. Selling agents respond with either a binding "yes" or "no".

Given this protocol, Kasbah provides buyers with one of three negotiation "strategies": anxious, cool-headed, and frugal – corresponding to a linear, quadratic, or exponential function respectively for increasing its bid for a product over time. The simplicity of these negotiation heuristics makes it intuitive for users to understand what their agents are doing in the marketplace.² This was important for user acceptance as observed in a recent Media Lab experiment [18]. A

2. Unlike other multi-agent marketplaces [29], Kasbah does not concern itself with optimal strategies or convergence properties. Rather, Kasbah provides more descriptive strategies that model typical haggling behavior found in classified ad markets.

larger Kasbah experiment is now underway at MIT allowing students to transact books and music [19].

Tete-a-Tete [30, 31] provides a unique negotiation approach to retail sales. Unlike most other on-line negotiation systems which competitively negotiate over price, Tete-a-Tete agents cooperatively negotiate across multiple terms of a transaction – e.g., warranties, delivery times, service contracts, return policies, loan options, gift services, and other merchant value-added services. Like Kasbah, this negotiation takes the form of multi-agent, bilateral bargaining but not using simple raise or decay functions as in Kasbah. Instead, Tete-a-Tete’s shopping agents follow an argumentative style of negotiation with sales agents (similar to [32]) and use the evaluation constraints captured during the Product Brokering and Merchant Brokering stages as dimensions of a multi-attribute utility (discussed in section 3.3). This utility is used by a consumer’s shopping agent to rank order merchant offerings based on how well they satisfy the consumer’s preferences. In essence, Tete-a-Tete integrates all three of the Product Brokering, Merchant Brokering, and Negotiation CBB stages.

3. Agent Technologies for Electronic Commerce

Most of the technologies supporting today’s agent-mediated electronic commerce systems stem from Artificial Intelligence (AI) research. From extracting meaning from ambiguous Web pages [17, 33] to planning trips to Hawaii [34, 35], from learning users’ music preferences [14, 13] to negotiating delivery contracts [36, 37] and deciding on which car to buy [38, 39], AI technologies will continue to provide software agents with increased know-how to successfully mediate electronic commerce transactions.

In this section, we review several AI technologies that support the systems described in section 2, discuss user interface challenges, then focus on issues and technologies concerning the next-generation agent-mediated electronic commerce infrastructure.

3.1. Recommender Systems

The majority of product recommender systems are developed using content-based, collaborative-based or constraint-based filtering methods as their underlying technology.

In *content-based filtering* [40, 41, 42, 43] the system processes information from various sources and tries to extract useful features and elements about its content. The techniques used in content-based filtering

can vary greatly in complexity. Keyword-based search is one of the simplest techniques that involves matching different combinations of keywords (sometimes in boolean form). A more advanced form of filtering is the one based on extracting semantic information from a document’s contents. This can be achieved by using techniques like associative networks of keywords in a sentence or price list, or directed graphs of words that form sentences.

Systems like BargainFinder and Jango try to collect information (e.g., product descriptions, prices, reviews, etc.) from many different Web information sources. These sources were intended to be read by humans and their content is rendered accordingly (i.e., in HTML). Different sources have different inputs (e.g., CGI-scripts, Java applets) and presentation methods, so recommender systems have to adjust their interaction methods depending upon the Web site. Since there is no standard way of defining and accessing merchant offerings, most recommender systems employ “wrappers” to transform the information from a specific Web site into a locally common format.

Different systems follow alternate approaches to creating wrappers. In BargainFinder, the Internet locations of on-line CD stores and the methods to access them (i.e., searching for a product and getting its price) are hand-coded by Andersen Consulting programmers. This method worked well at the beginning but is very hard to scale since it involves maintaining the wrapper for each site whenever it changes its access methods or catalog presentation format. Jango helps automate the creation of wrappers for new sites by generalizing from example query responses to on-line merchant databases. This technique is not perfect, but boasts a nearly 50% success rate in navigating random Internet resources [44].

Firefly uses a *collaborative-based filtering* technology [14, 45] to recommend products to consumers. Systems using collaborative techniques use feedback and ratings from different consumers to filter out irrelevant information. These systems do not attempt to analyze or “understand” the features or the descriptions of the products. Rather, they use consumers’ rankings to create a “likability” index for each product. This index is not global, but is statistically computed for each user on the fly by using the profiles of other users with similar interests. Products that are liked by similar-minded people will have priority over products that are disliked.

As in content-based approaches, *constraint-based filtering* uses features of items to determine their relevance. However, unlike most feature-based techniques which access data in their native formats, constraint-based techniques require that the problem and solu-

tion space be formulated in terms of variables, domains, and constraints. Once formulated in this way, however, a number of general purpose (and powerful) *constraint satisfaction problem* (CSP) techniques can be employed to find a solution [46, 47].

Many problems can be formulated as a CSP such as scheduling, planning, configuration, and machine vision problems. In PersonaLogic, CSP techniques are used in the Product Brokering CBB stage to evaluate product alternatives. Given a set of constraints on product features, PersonaLogic filters products that don't meet the given "hard" constraints and prioritizes the remaining products using "soft" constraints (which need not be completely satisfied).

Tete-a-Tete uses CSP techniques to assist shoppers in the Product Brokering, Merchant Brokering, and Negotiation CBB stages. This is achieved by consumers providing product constraints (as in PersonaLogic) as well as merchant constraints such as price, delivery time, warranty, etc. Hard and soft constraints are used to filter and prioritize products and merchants as well as construct a multi-attribute utility that is used to negotiate with the merchants. Tete-a-Tete's argumentative style of negotiation resembles a distributed CSP [48] with merchants providing counter-proposals to each customer's critiques [32].

3.2. User Interface Approaches

Traditional shopping experiences vary depending upon the needs of the consumer and nature of the product offerings. For instance, sometimes a shopper is just browsing without a specific intention to buy or sometimes the shopper intends to buy but is unfamiliar with the features of the specific product category (e.g., "I just need a camcorder whose tapes are compatible with my VCR."). Other times, the shopper intends to buy and has a deep understanding of the product category (e.g., "I need a S-VHS camcorder with x16 optical zoom.") Matching the system's user interface with the consumer's manner of shopping will likely result in greater customer satisfaction.

The user interface that most systems offer today is an "electronic catalog" which resembles an enhanced price list with search capabilities. Unfortunately, these searchable lists still make it hard for consumers to associate a product with their specific needs and afford less engaging shopping experiences than their physical-store counterparts.

One approach to help overcome these problems is the on-line mimicking of familiar physical-world shopping elements. For example, 3D VRML shopping malls have been developed to provide a more familiar shopping context. Although promising [49], these

shopping environments have not yet lived up to their expectations due to the awkwardness of navigating 3D worlds with 2D interfaces and other technical limitations (e.g., bandwidth).

Another example is the introduction of sales agent avatars – semi-animated graphical characters that interact in natural language with the consumer and feature a long-term consistent "personality" that remembers each customer, his or her shopping habits, etc. Anthropomorphized avatars (e.g., from Extempo [50]) attempt to mimic real-world sales agents to provide a more engaging on-line shopping experience and assist customers in finding the products that best meet their needs. Through immediate positive feedback and personalized attention, anthropomorphized sales agents can help build engaging, trusted relationships with customers [51]. However, the AI technologies behind the graphical representations of today's avatars are not yet up to meeting their users' expectations. Due to this and other reasons, the anthropomorphization of agents is still a controversial approach [52].

The issue of trust is very important in any agent system, especially when money is involved. A crucial issue in developing trust in agent systems is the ability of an agent to exhibit somewhat predictable behavior and to provide an explanation for its actions. For instance, a consumer can follow the decision process of a constraint satisfaction system like PersonaLogic much easier than that of a collaborative filtering system like Firefly which bases its recommendations on "invisible" clusters of like-minded people. In regards to complexity and predictability of behaviors, preliminary experiments with the Kasbah system showed that consumers preferred simple, predictable agents with pre-determined negotiation strategies over "smarter" agents that continuously adapted their behavior depending on an analysis of the marketplace.

It is safe to assert that, as with any software system, agents that mediate electronic commerce transactions can greatly benefit from well-designed and well-tested user interfaces [53].

3.3. Negotiation Mechanisms

Negotiation is a form of decision-making where two or more parties jointly search a space of possible solutions with the goal of reaching a consensus [37]. Economics and game theory describe such interactions in terms of protocols and strategies. The *protocols* of a negotiation comprise the rules (i.e., the valid actions) of the game. An example of a simple negotiation protocol is the Dutch auction where the only legal bidding action is an open outcry of "mine!" (or comparable) as

an auctioneer decrements the price of the good. We discussed other negotiation protocols in section 2.4.

For a given protocol, a bidder uses a rational *strategy* (i.e., a plan of action) to maximize his or her utility. Decision analysis tools help identify optimal strategies given a bidder's preferences and knowledge (e.g., motivation, valuation, risk, information asymmetry, etc.) [39].

Whereas economics research often focuses on partial and general equilibrium aspects of market-based negotiation, game theory research tends to focus on identifying optimal (self-interested) strategies and predicting outcomes for a variety of negotiation protocols [37]. A key idea from both of these research areas is that the specification of the protocol will have substantial, rippling effects on the nature of the overall system [37]. In other words, protocol design in the CBB Negotiation stage of agent-mediated electronic commerce should be considered carefully.

The research area that merges negotiation with software agents is the broad field of Multi-Agent Systems (MAS) which finds its roots in Distributed Artificial Intelligence (DAI). Early DAI work modeled negotiation as Distributed Problem Solving and assumed a high degree of *cooperation* among agents in order to jointly achieve a common goal [54, 55]. For example, the Persuader intermediary system combined multi-attribute utility theory and case-based reasoning to identify a mutually optimal deal during labor relation negotiations [56, 57]. More recent MAS work in market-based systems (e.g., AuctionBot [23]), on the other hand, assumes total self-interest and a high degree of *competition* among agents during negotiations for limited resources [58].

Much of the work in agent-mediated negotiations can be traced back to the *Contract Net* [59]. The original Contract Net was a distributed problem solving system designed for opportunistic, adaptive task allocation with agents announcing tasks, placing bids, and awarding contracts. Limitations of the original Contract Net Protocol (CNP) have been addressed in more recent work by Sandholm and Lesser [36, 58]. Related work includes Malone, et al.'s Enterprise system which allocates computer tasks using negotiation mechanisms [60] and protocols for automated coalition formation among agents [61, 62, 63, 64]. These latter protocols allow self-interested agents to cooperate on tasks (e.g., leverage economies of scale) without a priori relationships among their owners.

3.4. Infrastructure, Languages, Protocols

As discussed, there are already many agent-mediated electronic commerce systems, each roughly

focused on only one or two CBB stages (see Table 1). Ideally, we would be able to mix and match systems playing in complementary stages to provide a full consumer shopping experience. Unfortunately, these systems were not designed to interoperate in this way and linking these disparate systems together would require a good deal of work.

In fact, several of the systems discussed (e.g., BargainFinder and Jango), require proprietary "wrapper" techniques to "scrape" Web pages for product and merchant content. This is because Web pages are currently written in HTML (hypertext markup language) which is a data *format* language. In contrast, XML (extensible markup language) is a data *content* meta-language allowing for the semantic tagging of data [65, 66, 67]. Microsoft and Netscape have each promised support for XML with style sheets in their respective Web browsers to help replace HTML with XML as the language of the Web. The World Wide Web Consortium has recently proposed the first version of the XML specification.

However, XML is not a panacea for system interoperability. Even with tagged data, tags need to be semantically consistent across merchant boundaries at least for the full value chain of a given industry. CommerceNet and member organizations are working towards such common ontologies [68]. However, it's still an open question how transactional terms should be universally defined and who should manage their evolution.

Related agent-based languages and protocols include KIF (Knowledge Interchange Format) [69], KQML (Knowledge Query Manipulation Language) [70, 71], and Ontolingua [72], an ontology sharing protocol. These were designed so heterogeneous agents and systems could describe knowledge and communicate it meaningfully in order to interoperate. In electronic commerce, this knowledge would include the definitions and semantics of consumer profiles, merchants, goods, services, value-added services, and negotiation protocols (among others).

For business-to-business electronic commerce, the dominant protocol is EDI (Electronic Data Interchange). EDI is a set of ANSI and U.N. standard protocols for inter-business transactions [73, 74]. EDI facilitates large-scale, repetitive, pre-arranged transactions between businesses in specific industries with each industry adapting the EDI protocol to its specific needs. Standard EDI transactions are performed through expensive, proprietary Value-Added Networks (VANs). Although a pioneering protocol for inter-business electronic commerce, EDI has several disadvantages: it is ambiguous, expensive to implement and maintain, and it is focused on large scale

business-to-business transactions leaving small and medium-sized enterprises (SME) without a business-to-business transaction protocol standard. This forces business relationships to be established a priori and provides a disincentive to dynamically explore more lucrative deals.

Other electronic commerce protocol proposals that agents may need to “speak” include Internet-based EDI (EDIINT) [75], XML/EDI (a grassroots effort) [76], Information Content & Exchange (ICE) [77] for the exchange of on-line assets among companies, Open Buying on the Internet (OBI) [78] for high-volume, low-dollar business-to-business purchases, as well as a host of niche protocols such as Open Financial Exchange (OFX) [79] for financial transactions, Secure Electronic Transactions (SET) [80] for credit card transactions, and Open Profiling Standard (OPS) and Personal Privacy Preferences Project (P3P) [81] for defining privacy options for consumer profile data – to name only a few.

In addition to document and protocol standards, there is a need for electronic commerce component standards for objects and agents. There are several competing technologies in this space including the Object Management Group’s CORBA/IIOP (Common Object Request Broker Architecture/Internet Inter-ORB Protocol) [82, 83], Microsoft’s COM and DCOM [84], and Sun’s Java and RMI (Remote Method Invocation) [85, 86] as well as several mobile agent platforms such as ObjectSpace’s Voyager [87], Mitsubishi’s Concordia [88], General Magic’s Odyssey [89], and IBM’s Aglets [90] – several of which have been proposed for OMG’s Mobile Agent Facility (MAF) [91]. Requirements for open, heterogeneous component-based commerce systems include backward-compatibility to “legacy” systems, fault-tolerance, efficient performance, extensibility, scalability, security, some concurrency control, and some registry mechanisms to tie all of the pieces together. Many of these issues are core to multi-agent systems research, distributed database research, distributed systems research, and group communications research.

4. Conclusion & Future Directions

Today’s first-generation agent-mediated electronic commerce systems are already creating new markets (e.g., low-cost consumer-to-consumer and refurbished goods) and beginning to reduce transaction costs in a variety of business tasks. However, we still have a long way to go before software agents transform how businesses conduct business. This change will occur as Software Agent technologies mature to better manage ambiguous content, personalized preferences, com-

plex goals, changing environments, and disconnected parties. The greatest changes may occur, however, once standards are adopted and evolved to unambiguously and universally define goods and services, consumer and merchant profiles, value-added services, secure payment mechanisms, inter-business electronic forms, etc.

During this next-generation of agent-mediated electronic commerce, agents will enhance customer satisfaction and streamline business-to-business transactions, reducing transaction costs at every stage of the supply chain. At some critical threshold, new types of transactions will emerge in the form of dynamic relationships among previously unknown parties. Agents will strategically form and reform coalitions to bid on contracts and leverage economies of scale – in essence, creating dynamic business partnerships that exist only as long as necessary. It is in this third-generation of agent-mediated electronic commerce where companies will be at their most agile and markets will approach perfect efficiency.

5. References

- [1] P. Maes. “Agents that Reduce Work and Information Overload.” *Communications of the ACM*, vol. 37, no.7, pp. 31-40, 146, ACM Press, July 1994.
- [2] J. Bailey and Y. Bakos. “An Exploratory Study of the Emerging Role of Electronic Intermediaries.” *International Journal of Electronic Commerce*, vol. 1, no. 3, Spring 1997.
- [3] D. Hawkins, K. Coney, and R. Best. *Consumer Behavior: Implications for Marketing Strategy*. Business Publications, Inc., 1980.
- [4] K. Runyon and D. Stewart. *Consumer Behavior*, 3rd ed. Merrill Publishing Company, 1987.
- [5] F. Nicosia. *Consumer Decision Processes: Marketing and Advertising Implications*. Prentice Hall, 1966.
- [6] J. Howard and J. Sheth. *The Theory of Buyer Behavior*. John Wiley and Sons, 1969.
- [7] J. Engel and R. Blackwell. *Consumer Behavior*, 4th ed. CBS College Publishing, 1982.
- [8] J. Bettman. An Information Processing Theory to Consumer Choice. Addison-Wesley, 1979.
- [9] A. Andreasen. “Attitudes and Customer Behavior: A Decision Model.” In L. Preston (ed.), *New Research in Marketing*. California Institute of Business and Economics Research, University of California, 1965.
- [10] I. Terpsidis, A. Moukas, B. Pergioudakis, G. Doukidis, and P. Maes. “The Potential of Electronic Commerce in Re-Engineering Consumer-Retail Relationships through Intelligent Agents.” J.-Y. Roger, B. Stanford-Smith, and P. Kidd. (eds.), *Advances in Information Technologies: The Business Challenge*, IOS Press, 1997.
- [11] Y. Bakos. “Reducing Buyer Search Costs: Implications for Electronic Marketplaces,” *Management Science*, vol. 43, no. 12, December 1997.
- [12] PersonaLogic URL: <<http://www.personalogic.com/>>
- [13] Firefly URL: <<http://www.firefly.com/>>
- [14] U. Shardanand and P. Maes. “Social Information Filtering: Algorithms for Automating ‘Word of Mouth’.” *Proceedings of the Computer-Human Interaction Conference (CHI’95)*, Denver, CO, May 1995.
- [15] BargainFinder URL: <<http://bf.cstar.ac.com/bf/>>
- [16] Jango URL: <<http://www.jango.com/>>
- [17] R. Doorenbos, O. Etzioni, and D. Weld. “A Scalable

- Comparison-Shopping Agent for the World Wide Web.” *Proceedings of the First International Conference on Autonomous Agents (Agents’97)*. Marina del Rey, CA, February 1997.
- [18] A. Chavez, D. Dreilinger, R. Guttman, and P. Maes. “A Real-Life Experiment in Creating an Agent Marketplace.” *Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM’97)*. London, UK, April 1997.
- [19] Kasbah URL: <<http://kasbah.media.mit.edu/>>
- [20] O. Williamson. *Markets and Hierarchies: Analysis and Antitrust Implications*. The Free Press, New York, 1975.
- [21] OnSale URL: <<http://www.onsale.com/>>
- [22] AuctionWeb URL: <<http://www.ebay.com/aw>>
- [23] P. Wurman, M. Wellman, and W. Walsh. “The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents.” To appear in *Proceedings of the Second International Conference on Autonomous Agents (Agents’98)*, May 1998.
- [24] AuctionBot URL: <<http://auction.eecs.umich.edu/>>
- [25] J. Rodriguez, P. Noriega, C. Sierra, and J. Padget. “FM96.5: A Java-based Electronic Auction House.” *Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM’97)*. London, UK, April 1997.
- [26] Fishmarket URL: <<http://www.iiia.csic.es/Projects/fishmarket/newindex.html>>
- [27] J. Rodriguez, F. Martin, P. Noriega, P. Garcia, and C. Sierra. “Competitive Scenarios for Heterogeneous Trading Agents.” To appear in *Proceedings of the Second International Conference on Autonomous Agents (Agents’98)*.
- [28] R. Axelrod. *The Evolution of Cooperation*. Harper Collins, 1984.
- [29] C. Sierra, P. Faratin, and N. Jennings. “A Service-Oriented Negotiation Model Between Autonomous Agents.” *Proceedings of the Eighth European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW’97)*. Ronneby, Sweden, May 1997.
- [30] R. Guttman and P. Maes. “Agent-mediated Integrative Negotiation for Retail Electronic Commerce.” To appear in the *Proceedings of the Workshop on Agent Mediated Electronic Trading (AMET’98)*, Minneapolis, Minnesota, April 9, 1998.
- [31] Tete-a-Tete URL: <<http://ecommerce.media.mit.edu/Tete-a-Tete/>>
- [32] S. Parsons, C. Sierra, and N. Jennings. “Agents that Reason and Negotiate by Arguing.” To appear in the *Journal of Logic and Computation*.
- [33] J. Allen. *Natural Language Understanding*. Benjamin Cummings, 1987.
- [34] G. Linden, S. Hanks, N. Lesh. “Interactive Assessment of User Preference Models: The Automated Travel Assistant.” *Proceedings of the Sixth International Conference on User Modeling*, 1997.
- [35] D. Weld. “Planning-Based Control of Software Agents.” *The Third International Conference on Artificial Intelligence Planning Systems (AIPS’96)*. Edinburgh, Scotland, May 1996.
- [36] T. Sandholm. “An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations.” *Proceedings of the Eleventh National Conference on Artificial Intelligence (AAAI’93)*, Washington, D.C., 1993.
- [37] J. Rosenschein and G. Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers*. MIT Press, 1994.
- [38] P. Gmytrasiewicz, E. Durfee, and D. Wehe. “A Decision-Theoretic Approach to Coordinating Multiagent Interactions.” In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, August 1991.
- [39] R. Keeney and H. Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-offs*. John Wiley and Sons, 1976.
- [40] G. Salton. “The SMART Retrieval System.” *Experiments in Automatic Document Processing*. 1971.
- [41] N. Belkin and B. Croft. “Information Filtering and Information Retrieval.” *Communications of the ACM*, 35, No. 12:29–37, 1992
- [42] A. Moukas. “Amalthaea: Information Filtering and Discovery using a Multiagent Evolving System.” *Journal of Applied AI*, Vol. 11, No. 5, 1997
- [43] H. Lieberman. “Letizia: An Agent that Assists Web Browsing.” In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI’95)*. Montreal, Canada, 1995.
- [44] N. Kushmerick, D. Weld, R. Doorenbos. “Wrapper Induction for Information Extraction.” In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI’97)*, 1997
- [45] P. Resnik, N. Iacovou, M. Sushak, P. Bergstrom, and J. Riedl. “GroupLens: An Open Architecture for Collaborative Filtering of Netnews.” In *Proceedings of Computer Supported Cooperative Work (CSCW’94)*, 1994
- [46] E. Tsang. *Foundations of Constraint Satisfaction*. Academic Press, 1993.
- [47] V. Kumar. “Algorithms for Constraint Satisfaction Problems: A Survey.” *AI Magazine*, 13(1):32-44, 1992.
- [48] Z. Collin, R. Dechter, and S. Kaiz. “On the Feasibility of Distributed Constraint Satisfaction.” *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI’91)*, 1991.
- [49] A. Richmond. “Enticing Online Shoppers to Buy – A Human Behavior Study.” In the *Proceedings of the Fifth International World Wide Web Conference*. Paris, France, May 6-10, 1996.
- [50] Extempo URL: <<http://www.extempo.com/>>
- [51] J. Sculley. “The Knowledge Navigator.” Video, 1987.
- [52] A. Wexelblat and P. Maes. “Experiments on Anthropomorphizing Agents.” Submitted to *Interactions*, ACM Press.
- [53] B. Shneiderman and P. Maes. “Debate: Direct Manipulation vs. Interface Agents.” *Interactions: New Visions for Human-Computer Interaction*, vol. iv.6, December 1997.
- [54] S. Lander and V. Lesser. “Understanding the Role of Negotiation in Distributed Search Among Heterogeneous Agents.” In the *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*. Chambéry, France, August 1993.
- [55] M. Yokoo, E. Durfee, T. Ishida, and K. Kuwabara. “Distributed Constraint Satisfaction for Formalizing Distributed Problem Solving.” In the *Proceedings of the 12th IEEE International Conference on Distributed Computing Systems*, 1992.
- [56] K. Sycara. “The PERSUADER.” In *The Encyclopedia of Artificial Intelligence*. D. Shapiro (ed.), John Wiley and Sons, January, 1992.
- [57] Persuader URL: <<http://almond.srv.cs.cmu.edu/afs/cs/user/katia/www/persuader.html>>
- [58] T. Sandholm and V. Lesser. “Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework.” In the *Proceedings of the First International Conference on Multiagent Systems (ICMAS’95)*, San Francisco, 1995.
- [59] R. Smith. “The Contract Net Protocol: A High Level Communications and Control in a Distributed Problem Solver.” *IEEE Transactions on Computers*, 1980.
- [60] T. Malone, R. Fikes, K. Grant, and M. Howard. “Enterprise: A Market-like Task Scheduler for Distributed Computing Environments.” In *The Ecology of Computation*. B.A. Huberman (ed.), North-Holland Publishing Company, 1988.
- [61] G. Zlotkin and J. Rosenschein. “Coalition, Cryptography, and Stability: Mechanisms for Coalition Formation in Task Oriented Domains.” In *Proceedings of the National Conference on Artificial Intelligence*, Seattle, Washington, July 1994. [Preliminary version of this paper appeared in AAAI Spring Symposium. Also appeared as a MIT Center for Coordination Science Technical Report #168 and MIT Sloan School WP #3703-94, July 1994.]
- [62] O. Shehory and S. Kraus. “Coalition Formation among Autonomous Agents: Strategies and Complexity.” C. Castelfranchi and J. Muller, (eds.), *From Reactions to*

- Cognition: Proceedings of MAAMAW'93*. Springer Verlag Publishers, Lecture Notes in Artificial Intelligence, 957, 1995.
- [63] S. Ketchpel. "Coalition Formation Among Autonomous Agents." C. Castelfranchi and J. Muller, (eds.), *From Reactions to Cognition: Proceedings of MAAMAW'93*. Springer Verlag Publishers, Lecture Notes in Artificial Intelligence, 957, 1995.
- [64] T. Sandholm and V. Lesser. "Coalition Formation among Bounded Rational Agents." *14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, Montreal, Canada, 1995.
- [65] J. Bosak. "XML, Java, and the Future of the Web." March 1997.
- [66] "Extensible Markup Language (XML)." World Wide Web Consortium (W3C) Working Draft, 17 November 1997.
- [67] Summer Institute of Linguistics XML URL: <<http://www.sil.org/sgml/xml.html>>
- [68] J. Tenenbaum, T. Chowdhry, K. Hughes. "eCo System: CommerceNet's Architectural Framework for Internet Commerce White Paper & Prospectus." Second draft, April 1997. <<http://www.commerce.net/eco/index.html>>
- [69] Knowledge Interchange Format (KIF) URL: <<http://logic.stanford.edu/kif/>>
- [70] T. Finin, D. McKay, R. Fritzson, and R. McEntire. "KQML: An Information and Knowledge Exchange Protocol." K. Fuchi and T. Yokoi (ed.), *Knowledge Building and Knowledge Sharing*, Ohmsha and IOS Press, 1994.
- [71] T. Finin, Y. Labrou, and J. Mayfield. "KQML as an Agent Communication Language." J. Bradshaw (ed.), *Software Agents*, MIT Press, 1997.
- [72] T. Gruber. "Ontolingua: A Mechanism to Support Portable Ontologies." Stanford University, Knowledge Systems Laboratory, Technical Report KSL-91-66, March 1992.
- [73] K. Blacker. *The Basics of Electronic Data Interchange, 2nd Edition*. D R Lambert, February 1994.
- [74] Premenos' Electronic Commerce Resource Guide to ANSI X12 and UN/EDIFACT Standards URL: <<http://www.premenos.com/standards/>>
- [75] R. Drummond, M. Jansson, and C. Shih. "Requirements for Inter-operable Internet EDI." Internet Engineering Task Force (IETF) Working Draft, 4 December 1997. <<ftp://ietf.org/internet-drafts/draft-ietf-ediint-req-04.txt>>
- [76] B. Marchal, N. Mikula, B. Peat and D. Webber. "Guidelines for using XML for Electronic Data Interchange, Version 0.03." M. Bryan (ed.), 21 October 1997. <<http://www.geocities.com/WallStreet/Floor/5815/guide.htm>>
- [77] ICE Press Release: <http://www.firefly.net/company/pr_2-1-98_ICE.fly>
- [78] "Open Buying on the Internet: A Standard for Business-to-Business Internet Commerce." SupplyWorks white paper. <<http://www.supplyworks.com/obi/white-paper.html>>
- [79] Open Financial Exchange (OFX) URL: <<http://www.ofx.net>>
- [80] Visa's Secure Electronic Transactions (SET) URL: <<http://www.visa.com/cgi-bin/vee/nt/ecom/aset/main.html>>
- [81] Personal Privacy Preferences (P3) Project Architecture Working Group. "General Overview of the P3P Architecture." World Wide Web Consortium (W3C) Working Draft, 22 October 1997. <<http://www.w3.org/TR/WD-P3P-arch.html>>
- [82] S. Vinowski. "CORBA: Integrating Diverse Applications Within Distributed Heterogeneous Environments." *IEEE Communications Magazine*, Vol. 14, No. 2, February 1997.
- [83] CORBA References: <<http://www.cs.wustl.edu/~schmidt/corba-papers.html>>
- [84] D. Rogerson. *Inside COM: Microsoft's Component Object Model*. Microsoft Press, 1997.
- [85] P. Sridharan. *Advanced Java Networking*. Prentice Hall, 1997.
- [86] "RMI and Java Distributed Computing." SunSoft white paper, 1997. <<http://www.javasoft.com/features/1997/nov/rmi.html>>
- [87] ObjectSpace's Voyager URL: <<http://www.objectspace.com/voyager/>>
- [88] Mitsubishi's Concordia URL: <<http://www.meitca.com/HSL/Projects/Concordia/>>
- [89] General Magic's Odyssey URL: <http://www.genmagic.com/html/more_odyssey.html>
- [90] IBM's Aglets URL: <<http://www.trl.ibm.co.jp/aglets/>>
- [91] Object Management Group's (OMG) Mobile Agent Facility (MAF) URL: <http://www.omg.org/library/schedule/CF_RFP3.htm>