Brokering custom-built information products: The IBIZA project

Ashish Arora[♦], Greg Cooper[♠], Ramayya Krishnan[♦♣], Rema Padman[♦♠], Rudy Setiawan[•], Kevin Teo[•] [1]

Abstract

This paper describes IBIZA, a workbench for designing electronic markets for information products. IBIZA is a testbed for performing computational experiments about market-based approaches (e.g., auctions) to support negotiations between buyers and sellers. We outline the key technical and economic issues encountered in the design of IBIZA. We use examples from our work on designing a software agent-based electronic market for automated model development illustratively.

[1] ♦ The Heinz School of Public Policy and Management, Carnegie Mellon University
♣ The Institute for Complex Engineered Systems, Carnegie Mellon University
♠ The Center for Biomedical Informatics, University of Pittsburgh
• The Information Networking Institute, Carnegie Mellon University

1. Introduction

There is general agreement that "negotiation of terms" is among the least supported of the trading processes in electronic commerce (Kumar and Feldman, 1999, Beam et al., 1999). When the negotiation is about a single dimension such as price and when a seller (buyer) needs to negotiate simultaneously with multiple buyers (sellers), auctions have become the method of choice on the Net. They are relatively easy to implement and have been accepted by consumers (well known examples include the Yankee auctions of Onsale and the universal reverse auction of CXN) (Beam and Segev, 1998). Extending auction schemes to accommodate multi-dimensional negotiations is still at an early stage (Bichler, Kumar, Jhingran, 1998). In parallel, there has been recent work on peer to peer negotiations between software agents (e.g., the Kasbah project at MIT (Chavez and Maes, 1996)) as well as the application of auction-based schemes to solve difficult resource allocation problems (Wurman et al., 1998).

Our work on the IBIZA project is complementary to these works but distinct in its focus on markets featuring custom-built information products with software agents for market players. By custom-built information products, we mean an information product being created on demand to meet the specifications of a buyer. We have used a complex task of automated model development as an example of custom product development (see scenario below).

IBIZA is a workbench that facilitates the design and evaluation of electronic markets and the agents that participate in them. Specifically, it enables designers to instantiate and simulate the behavior of a market featuring a specific market mechanism (e.g., a vickrey auction) with agents (sellers) possessing specific assets (e.g., machine learning methods) using specialized sets of bidding strategies. The objective is to facilitate computational experimentation with and testing of hypotheses about market mechanisms and bidding strategies in electronic markets for information products and services.

The objective of this paper is to provide an overview of IBIZA. We first introduce a scenario that we have used as a prototypical example to drive the design of IBIZA. The

scenario describes a market for predictive model development services. Using this scenario to provide context, we discuss key technical and economic issues encountered in the design of IBIZA. We conclude with a discussion of experiments and technical issues that can be investigated with IBIZA.

2. Illustrative Scenario

The context*: A researcher wants to develop a predictive model to assist a clinician in treating community acquired pneumonia. She has the MCHD data set that contains information on inpatients discharged from 78 hospitals in 23 states between July 1987 and December 1988. It tracks over 250 pieces of clinical information referred to as key clinical findings (KCF). The KCF's include patient demographic characteristics, history, physical examination findings, lab and radiology results collected during up to 3 reviews during the hospitalization. The model should predict mortality of hospitalized patients from their findings at initial presentation with pneumonia. Such predictions would be useful to clinicians since they must decide about where to treat patients with pneumonia. Treating patients at home is less expensive than treating patients at the hospital, and patients with milder cases of the disease are likely to be more comfortable at home.*

The IBIZA market: The researcher connects to the model development broker in order to "buy" the model (i.e., more specifically, to have a custom model built using her data set). The broker runs a reverse market in which the sellers are software agents who have the assets (i.e., machine learning methods) to build the model (see Figure 1).

First, she prepares the custom data set to be used in model creation by selecting 75 out of the 200+ variables that she thinks are relevant to the learning task.

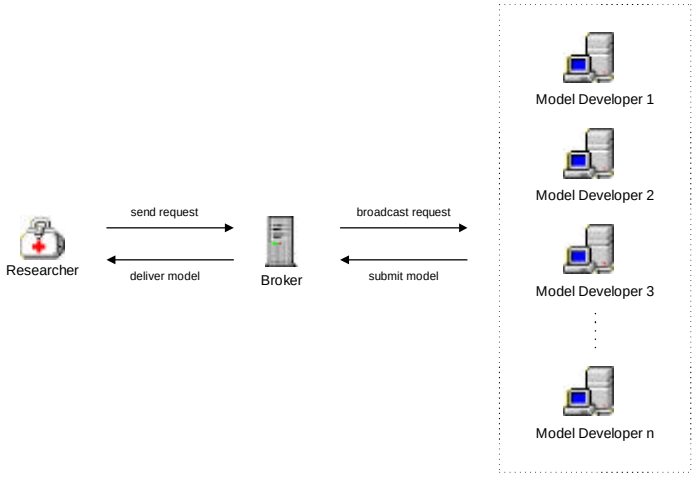

Figure 1: The IBIZA market architecture

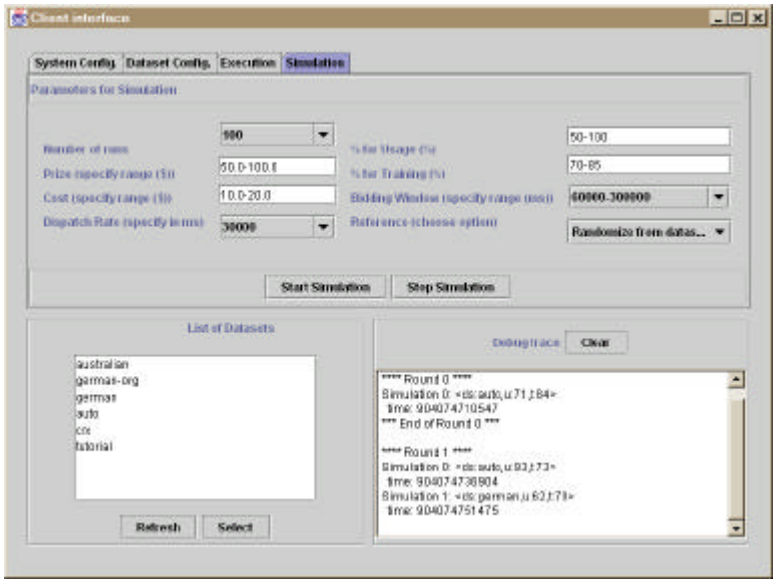

Figure 2: The Consumer Interface in IBIZA

She divides the 2000 records in her data set into two parts, a training data set consisting of 1200 records and an evaluation data set consisting of 800 records to test the quality of the models that will be developed in the market. Using the consumer interface applet (see Figure 2), she then poses a request to the broker consisting at the very minimum of the following.

- The training and evaluation data sets (URL's of these resources)
- Metadata about the data set indicating types and value ranges of variables
- The maximum price she is willing to pay
- The minimum accuracy she wants the model to provide
- A parameter trading off price to quality (accuracy)
- The time window (deadline) by when she wants the model development completed

The broker processes this request (see Figure 3) and executes the market mechanism which for the purpose of this illustration is a ***prize scheme***. In this scheme, the broker generates an object called an **ask** from the user request. The ask object declares the prize (usually equal to the price the user is willing to pay) and with the exception of the evaluation data set, supplies the other information provided by the user to the broker. Seller agents evaluate if the ask object is worth bidding on, and if it is, use their machine learning method to build the model using the training data. This custom built model is encapsulated as an applet and submitted to the broker as part of the **bid**. The bids (or more specifically, the model submitted with the bid) are evaluated by an evaluation service on the evaluation data withheld from the seller agents by the broker and ranked in terms of their predictive accuracy. The model that has highest quality is deemed to be

winner in the prize scheme. The transaction concludes with the seller agent that developed the winning model being compensated by the broker and the model being transmitted to the user.

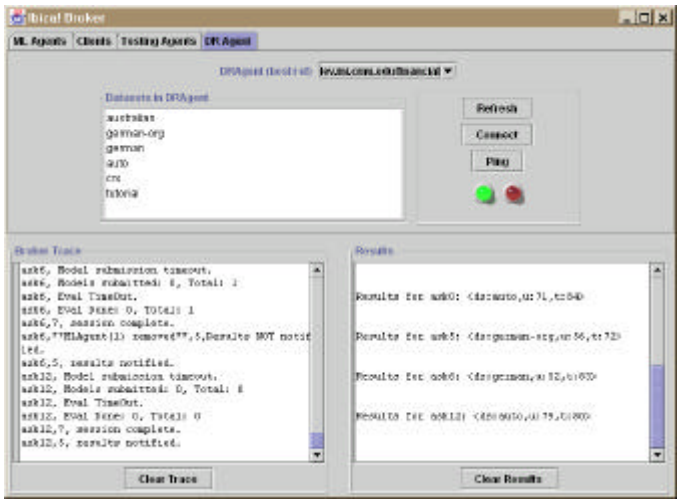

Figure 3: The IBIZA Broker Console

The rationale for the architecture and the use of market mechanisms for coordinating the negotiations between the buyers and sellers can be understood in terms of transaction costs, asset specificity and product complexity. The IBIZA environment assumes electronic transactions in information products with the following characteristics.

- While the cost of creating the asset – the learning algorithm – can be high, the marginal cost of producing the information product (i.e., the model) being transacted is comparatively low. Since production costs and communication costs are low, a market-based approach to acquiring these products is indicated (Malone, 1987).
- The processes used to create the information products have low asset specificity. For example, methods such as learning algorithms are not site specific in that they do not need specific computational platforms. Neither are they function specific in the sense that a learning algorithm could be applied to any learning task. Clearly, some methods are better suited to some tasks but the method is applicable nonetheless. Markets are indicated when the assets required to produce the product have low asset specificity.

- Evaluation of the product is not complex. In the case of models, their predictive accuracy can be tested in an objective manner using the evaluation data. Since a standardized method for evaluating products is available, markets can be used as the coordination mechanism between the buyers and sellers.

3. Analysis

In this section we present a brief overview of the key components of markets designed using IBIZA. These components can be understood in terms of the participants and the interactions between them. The interactions are determined by the market mechanism in use. The participants in IBIZA are:

- *Producers/Sellers:* These are software agents in IBIZA. They can be endowed with assets. In our context, each agent has a specific learning algorithm as an asset. This learning method is used to construct models from the training data sets submitted by the user.
- *Buyers:* These are human agents at present in IBIZA. However, there is no constraint on these being user software agents as well. In our context, buyers are users who would like models built to specification.
- *Broker:* The broker is the intermediary that provides match making services using one of several market mechanisms. At present, two schemes – a generalized auction procedure and a prize scheme are available for use in IBIZA. The choice of market mechanism and the type of seller agents determine the information processing requirements of the broker. This is discussed in more detail below.
- *Support service provider:* The performance of match making might require resources (e.g., CPU cycles and methods to evaluate and rank models submitted as part of the bid) that the broker does not possess. In IBIZA, model evaluation is an example of such a service. At present, model evaluation services are provided by a single provider and coordinated by the broker. However, alternatives are possible including secondary markets for this service as well as strategic relationships between certified model evaluators and sellers that do not involve coordinating with the broker.

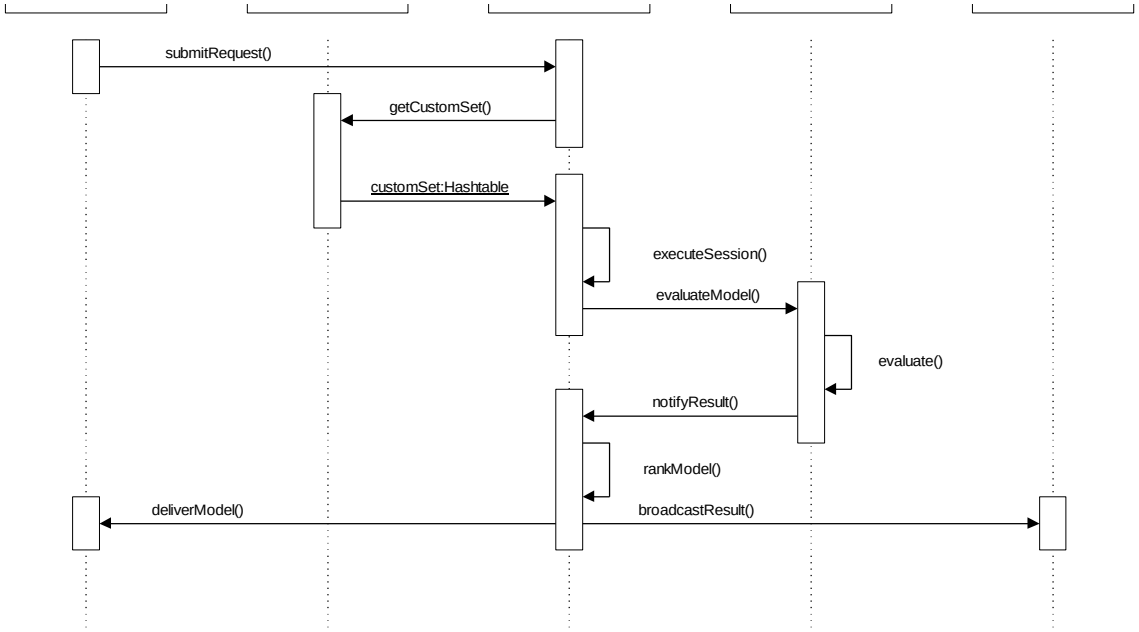

Figure 4: The IBIZA market mechanism time sequence diagram

- Market mechanisms: Two market mechanisms are available for use in IBIZA. They are a prize scheme and a generalized auction procedure. The prize scheme is a special case of the auction. The market mechanism determines the interactions between the players as well as the information processing requirements (see Figure 4). We briefly describe each of the schemes that are implemented in IBIZA.
  - The auction scheme in IBIZA is a two step procedure. The first step begins with the broker notifying all potential participants (it can be all sellers registered with the broker or just a subset of them that the broker selects based on reputation established over time) about the request received from the buyer. This is done using an **ask** object (see Figure 5), a container class used to package all the information required by seller agents to decide if they would like to bid.

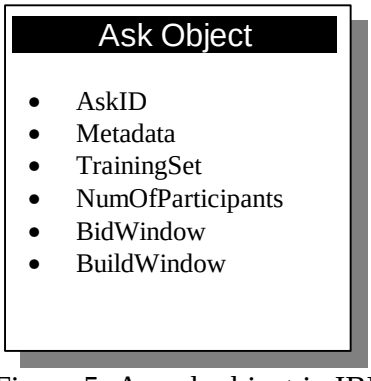

Figure 5: An ask object in IBIZA

Here we briefly elaborate on our discussion of the ask object in the previous section. The attribute *Numofparticpants* is set by the broker and specifies the number (not the

identity) of seller agents being invited to bid in a market session. This information is used by seller agents as an input to their bidding strategy. Experiments with seller bidding strategy may generate information requirements that impact on broker policy and the attributes of the broker controlled ask object is a good example of such an interaction and how it is managed in IBIZA. The BidWindow and BuildWindow are attributes that specify the time that is available to the seller agents to build their models. At present, agents in IBIZA do not have the ability to independently allocate the BidWindow (the total amount of time available to build and evaluate the model submitted as part of a bid). This is because we are currently experimenting with a market structure in which evaluation services are centrally coordinated by the broker. If the market structure for evaluation services were modified, this would permit seller agents to independently manage the bid window. Once again such as change would cause a change to the ask object. IBIZA provides the functionality to manage these changes in the market infrastructure by enabling designers to add components and instantiate available components in a parameterized manner.

Each seller agent processes the ask object and decides whether to participate in the auction. If they choose to do so, this is followed by the generation of a **bid** object (see Figure 6).

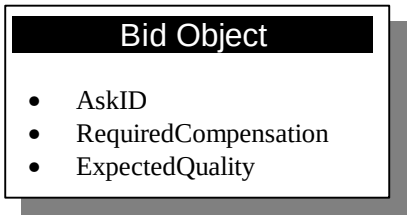

Figure 6: The IBIZA bid object

The bid object at a minimum indicates the ask that is being responded to and price quality vector that specifies the expected compensation and the promised quality. Investigation of alternative bidding strategies ranging from simple “bid on any ask” to more complex strategies based on expected gains from bidding is facilitated by allowing bidding strategy components to be plugged into IBIZA agents. Upon receiving all bids from the sellers, the broker chooses a subset of the bidders (can be one or more participants) that it will award the contract to develop the product. How

the broker should make this decision is an interesting research issue as well. One possibility is to award contracts to all bidders (in fact this defaults to the prize scheme discussed below). Other possibilities include maximizing welfare or net surplus (for details see Arora et al., 1999). At present, we have components implementing both of these alternative strategies in the broker and designers have a choice of using one strategy or the other in their computational experiments. Other strategies may be added in an evolutionary manner as well by encapsulating them into IBIZA compliant components. The sellers who have been awarded contracts submit the models they built to the broker for evaluation. These are evaluated and ranked. Currently, the ranking is limited to a two dimensional procedure based on the price and quality submitted as part of the bid. The relative weights on these factors is determined by the user as part of the user request. The model that best meets the buyers requirement is delivered to the buyer and the seller is compensated as per the contract issued by the broker at the end of the first step. Sellers that did not submit the winning bid get consolation awards or are charged a penalty if they do not meet the terms specified in their contract.

- The prize mechanism in IBIZA is a special case of the auction and works like a contest. It is a one step procedure that begins with all sellers registered with the broker being notified of the user request. The user has indicate the prize they would like to award as part of their request. Each seller can decide if they want to participate. If they do, they are required to build the product (i.e., the model) before bidding. This is equivalent to each bidder being awarded the contract to build the model in the auction procedure. The bids are ranked in terms of quality. The seller with the highest quality model wins the prize and the model is transmitted to the buyer. The prize scheme imposes fewer information processing requirements on the seller agents and the broker. The key decision to be made by the seller agents relates to their participation in the contest. Unlike the auction where the seller agents have to set a price quality vector, the only component of the bid here is the model. The broker has to rank the bids and award the prize to the winner and transmit the model to the

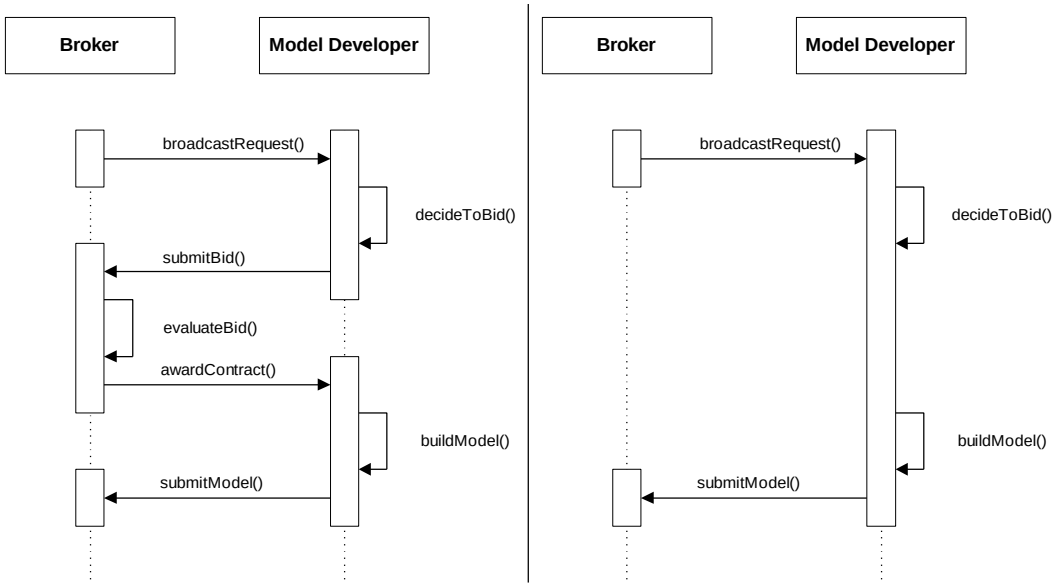

Figure 7: Difference between Auction and Prize schemes

buyer. No complex decisions need to be made by the broker about who to invite to bid or to whom to award contracts. The differences in the sequence of interactions between the auction scheme and the prize scheme is summarized in Figure 7.

4. Discussion

The availability of an environment such as IBIZA allows a range of technical and economic issues related to the use of market-oriented approaches for brokering information products to be investigated. Given our space limitations, we briefly outline a few issues we are investigating in ongoing research. Interested readers should refer to Arora et al. (1999).

Mechanism design: At present IBIZA offers two alternative schemes, a generalized auction procedure and a specialized variant of it called the prize scheme. We are currently conducting computational experiments to compare the two schemes. In particular, we are interested in the interaction between the effort that should be invested by the market maker (the broker) in the two step auction scheme to select sellers to award contracts and the quality of the outcomes such as Net Surplus. We are interested in how complex information intensive seller selection strategies compare with simple adaptive strategies that learn over time. While we are conducting our experiments using these two schemes, other market mechanisms could be easily added to IBIZA

Bidding Strategy of seller agents: At present, seller agents in IBIZA can be implemented using either a random bidding strategy or a decision tree-based approach that takes expected values of winning into account. We are interested in a comparative study of the interaction between bidding strategy and the market mechanism. A sample strategy matrix that can be used to guide experiments in this space is shown below. While there has been previous work in this area (Mcafee and McMillan, 1987; Milgrom, 1989, Oliver, 1996), the studies were not conducted in the IBIZA context of internet-based auctions or in the context of custom built information products with the characteristics described in Section 2.

Table 2: Bidding Strategy Matrix

Bidding Strategy	Mechanism Type	Decision Type
Naive	Auction Prize	Bid = Cost Plus x% Enter if expected cost less than threshold level
Adaptive	All mechanisms	Start with Naive, randomly vary strategy and select in response to feed back.
Rational	All mechanisms	Compute optimal bidding response given priors about the distribution of rivals' costs and capabilities.

Technology Issues: While the two examples we have discussed relate to design of markets and agents from an economic perspective, we have used IBIZA to also investigate technology infrastructure issues. Since the custom product being transacted in the examples we have used is a model, how should this model be created, submitted for evaluation and eventually delivered to the buyer? How should the model be communicated from one player to another. In addition to these being technology choices, they also have implications in relation to barriers to entry into the IBIZA market as well as to switching costs for agents choosing to leave the IBIZA-based market to participate in another market for information products. We compared four alternatives (Table 1) and chose the Java distributed object environment as the technology infrastructure for model shipment and delivery.

	TCP/IP Sockets	Remote Procedure Calls	Java Remote Method Invocation	CORBA IIOP
Platform Independence	Yes	No	Yes	Yes
Language Independence	Yes	No	No	Yes
Implementation Complexity	High	Low	Low	Low
Supports Object Serialization	No	No	Yes	Depends on implementation
3rd Party Package Necessary	No	No	No	Yes
Remote Object Execution	No	No	Yes	Yes

Table 1. Distributed Communication Technology Comparison.

The main reason for choosing Java over CORBA was the lightweight nature of the implementation and the fact that the Java distributed object model is available with the Java toolkit. Models are packaged as Java objects and communicated using object serialization capabilities available in Java. Of course, this does not preclude any seller agent from implementing its learning methods in languages such as C++. In experiments we have conducted, all our agents used learning methods implemented in C++. Only the products they create were encapsulated and transmitted as Java objects. Interoperability is ensured by requiring any model developed in the IBIZA markets to subscribe to a standard API allowing different players to evaluate and execute the model. Space limitations preclude a detailed discussion of this issue. The interested reader is referred to a detailed technical report (Arora et al, 1999).

Bibliography

1. "CXN, Inc." *http://www.cxn.com.*
2. "Onsale, Inc." *http://www.onsale.com*, Menlo Park, CA.
3. Arora, A., Cooper, G., Krishnan, R., Padman, R., Setiawan, R., Teo, K., The IBIZA project, Heinz School Working Paper, 1999.
3. Beam, C., and Segev, A. (1998). "Auctions on the Internet: a Field Study." *CITM Working Paper 98-WP-1032*, Haas School of Business, University of California at Berkeley, Berkeley, CA.
4. Beam, C., Bichler, M., Krishnan, R., Segev, A., "On Negotiations and Dealmaking in Electronic Markets", Information Systems Frontiers, forthcoming.

5. Bichler, M., Kumar, M., and Jhingran, A. (1998b). "Multi-attribute Competitive Bidding." *Working Paper*, IBM Thomas J Watson Research Center, Yorktown Heights, NY.
6. Chavez, A., and Maes, P. (1996). "Kasbah: An Agent Marketplace for Buying and Selling Goods." , MIT Media Lab, http://agents.www.media.mit.edu:80/groups/agents/Publications/kasbah-paam96.ps.
7. Kumar, M., and Feldman, S. (1999). "Internet Auctions." , IBM Institute for Advanced Commerce, Yorktown Heights, NY.
8. McAfee, R. P., and McMillan, J. (1987). "Auctions and Bidding." *Journal of Economic Literature*, 25(2), 699-738.
9. Milgrom, P. (1989). "Auctions and Bidding - a Primer." *Journal of Economic Perspectives*, 3(3), 3-22.
10. Oliver, J. R. (1996). "A Machine Learning Approach to Automated Negotiation and Prospects for Electronic Commerce." , University of Pennsylvania, http://opim.wharton.upenn.edu/~oliver27/papers/jmis.ps.
11. Wurman, P., Walsh, W., and Wellman, M. (1998). "Flexible Double Auctions for Electronic Commerce: Theory and Implementation." *Decision Support Systems*.
12. Zeng, D., and Sycara, K. (1996). "Bayesian Learning in Negotiation." *Working notes of the AAAI 1996 Stanford Spring Symposium Series on Adaptation, Co-evolution, and Learning in Multiagent Systems.*