An Evolutionary Framework for Studying Behaviors of Economic Agents

Wolfgang Ketter
Dept of Computer Science
and Engineering
University of Minnesota
ketter@cs.umn.edu

Alexander Babanov
Dept of Computer Science
and Engineering
Dept of Economics
University of Minnesota
babanov@cs.umn.edu

Maria Gini
Dept of Computer Science
and Engineering
University of Minnesota
gini@cs.umn.edu

ABSTRACT

We propose an evolutionary framework for studying agents that interact in electronic marketplaces. We describe how this framework could be used to study the dynamics of interaction and evolution of agent strategies. We present experimental results from a simulated market, where multiple service providers compete for customers using different pricing strategies. The results show that service providers having different strategies and capacities occupy different niches in the market.

Categories and Subject Descriptors

K.4.4 [**Computers and Society**]: E-commerce; I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence

General Terms

Social, Economics, Simulation

Keywords

Evolution, Adaptation and Learning

1. INTRODUCTION

Online marketplaces are gaining popularity among companies seeking to streamline their supply chains. For buyers such marketplaces can significantly ease the process of finding, comparing and coordinating providers, while for sellers marketplaces provide access to a much broader customer base [21].

Intelligent software agents can significantly reduce the burden of market exploration by sifting through the avalanche of information and performing bulky calculations to promptly provide a human decision maker with a refined list of alternatives. The sheer speed of automated negotiations and decision support can reduce the cost of committing to resources between submitting a bid and receiving a bid award or a rejection reply. In other words, agents can help human decision makers to make better choices faster.

In our research we are interested in studying how autonomous agents can maximize profits when requesting or submitting bids. Agents can use many different strategies to price their products when they submit bids (see [14] for an analysis of several pricing strategies).

We are specially interested in analyzing the performance of different bidding strategies in the context of the MAGNET (Multi-AGent NEgotiation Testbed) research project [8]. MAGNET agents participate in first-price, sealed-bid, reverse combinatorial auctions over collections of tasks with precedence relations and time constraints.

A major issue is assessing how good a strategy is. This is especially challenging when dealing with combinatorial auctions, since modeling analytically how a strategy will perform is very hard, if not impossible. Unfortunately, there are not enough real-world data available for a comprehensive testing of the effects of using different strategies.

In this paper we will mainly focus on the question: "How can we compare strategies used by agents, when not enough data are available?" Our proposed method is to design a large-scale test environment atop an evolutionary approach to economic simulation, and let the evolution of the market decide which strategies are most suited.

We start by proposing in Section 2 the use of an evolutionary approach to study the dynamics of interaction and evolution of strategies for agents that interact in a marketplace. In Section 3 we describe how to construct such an evolutionary system. In Section 4 we present a case study of a simulated market where multiple service providers compete for customers, and where profitability is the criterion used to stay in business. The experimental results we show conform to expectations. Service providers having different strategies and capacities occupy different niches in the market. Finally, in Section 5 we compare our approach with other methods.

2. WHY AN EVOLUTIONARY APPROACH?

We are interested in studying the performance of agents in a simulated marketplace, and develop an understanding of the properties of automated and mixed-initiative auction-based trading societies.

A major obstacle in the way of understanding the prop-

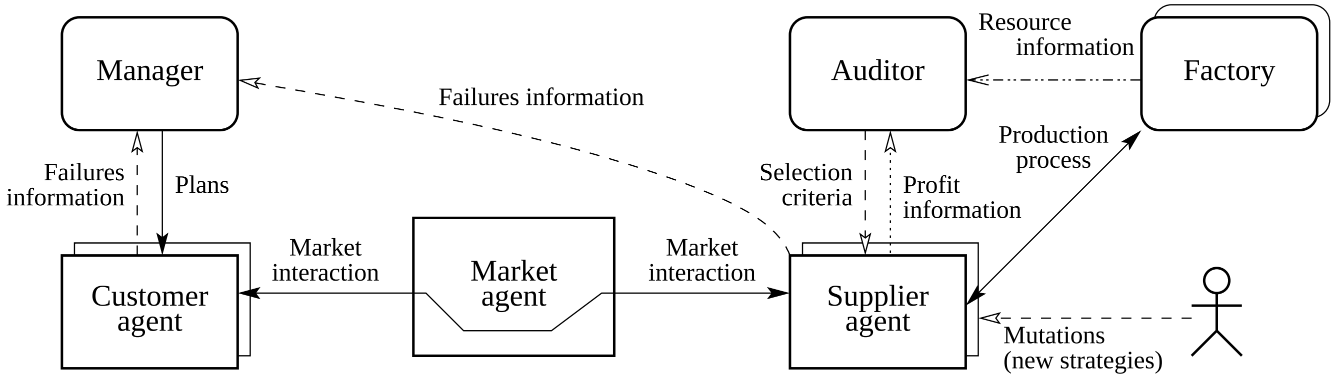

Figure 1: A multi-agent system architecture adjusted to the evolutionary paradigm. The rounded boxes show specific evolutionary components.

erties of multi-agent systems is the lack of tractable data. Publicly available data are scarce and insufficient for exhaustive testing, while private data sets are expensive and not always suitable for research purposes.

We propose a way of employing an evolutionary approach to economic simulation that will make up for the scarcity of data, while offering a scientific approach to data collection and a systematic tool for experimentation.

The methodology we use is based on the evolutionary approach to game theoretical problems. Evolutionary game theory [29] studies equilibria of games played by populations of players, where players are myopically rational and have conflicting interests.

In evolutionary systems there is no fitness function, instead there is a rule which governs survival of society members based on their success. So, the "fitness" of the players derives from the success each player has in playing the game governed by the natural selection.

In our case the players are customer and supplier agents, and their fitness is determined by the strategies they use to secure profit. Agents which do not perform well, because of their strategy, will eventually disappear from the market.

A similar approach has already been used successfully in other domains [5]. They tested different agent strategies in the context of the Nash's demand game. The strategies that survived are better, on average, than the ones who had to leave the market. Their simulation shows that the evolutionary approach performs well in selecting equilibria, even though the predicted population distributions are often different from the results obtained analytically.

The major benefits of using an evolutionary approach include:

- Any multi-agent system, such as MAGNET [8], is governed by a magnitude of parameters, many of which are continuous variables. The search space of the system is immense, thus rendering any systematic testing very hard if not impossible. Using an evolutionary approach gives us a way of searching the space of possible parameters more efficiently.
- A thorough study of agent strategies requires information about the behaviors of other agents in the system. The evolutionary approach solves this problem by enclosing all the agents in a self-sufficient system, where they can observe each other's behavior and influence each other's behaviors.
- The evolutionary approach allows for formation of complex spatio-temporal patterns of behavior that are not observable at the level of individual agents. Examples studied by other researchers range from the emergence of cooperation in an otherwise selfish society [4, 3] with possible formation of spatial patterns of strategic interaction [18] to ostracism and neighborhood effects [11] and natural phenomena, like fish schools [15].

3. ARCHITECTURAL DESIGN OF AN EVOLUTIONARY FRAMEWORK

We propose to design a large-scale test environment atop an evolutionary approach to economic simulation. The architecture of our proposed system, illustrated in Figure 1, is as follows:

The **Manager** generates and distributes plans to customer agents. It observes the rate at which suppliers fail to stay solvent in the market, and adjusts the frequency of arrival of customers to keep the rate of failures reasonably low, yet not zero. Having a rate of failures greater than zero puts some pressure on agents that use computationally overly intensive strategies.

The **Auditor** evaluates the performance of the supplier agents strategies based on the average profit over a specified period of simulation time. Agents that make negative profit are removed from the market. Whenever the average profit in the market exceeds some specified value, the auditor introduces a new supplier agent with a strategy that is chosen from the pool of all the strategies in the market, weighted by the number of suppliers that execute them. The auditor maintains a pool of "retired" strategies, i.e. strategies that were completely eliminated from the market, and eventually tries to put them back in the market. That allows retired strategies to try to take over the market in some more favorable time.

One instance of the **Factory** is assigned to each supplier agent to keep track of resource availability and existing commitments. The size and types of products produced in a factory are determined by the auditor upon creation of the corresponding supplier agent. Customer agents look for suppliers that can satisfy their needs.

Both **Customer agents** and **Supplier agents** make decisions in a completely autonomous way, without any human intervention. The supplier agent coordinates its resource commitments with its own factory.

Human participants can submit new strategies to the pool of possible **mutations**. "Mutant" strategies are introduced to the market after the market has reached a dynamic steady state.

The rationale behind our choice of an evolutionary framework is that it is able of providing results without requiring a complex theory of agent motivation, optimization criteria, or strategic interaction. The framework is determined by the motives of individual agents, the rules of agents' interaction, and the governing selection criterion. Given that, the evolutionary development of the system provides the dynamic information on the macroscopic behaviors of the society.

A brilliant explanation of the relation between micromotives of agents and macrobehavior can be found in [24]. Evolutionary frameworks have been used extensively in Economics [19, 22, 28].

3.1 Reproduction, Mutation and Introduction of New Strategies

One of the cornerstones of the evolutionary approach is the need for a large and diverse population of agents. A common solution to this issue is to describe agents' strategies in terms of gene sequences and to use cross-breeding and mutations to ensure the desired diversity.

Agents can employ a variety of strategies, such as Q-Learning, Neural Networks, Game Theoretic models, Genetic Algorithms and alike. It is hard to imagine that each and everyone of the strategies mentioned above can easily be encoded in a gene sequence. It is even harder, if not impossible, to maintain the compatibility between gene sequences of different strategies. In practice, it is pretty difficult to come up with an encoding for even well studied problems [10], let alone for complex domains like the MAGNET system.

We address the problem of reproduction and mutations by generalizing the concept of gene pool. We illustrate our approach by designing and investigating a simple model of a suppliers' and customers' community in Section 4.

Our proposed approach is to maintain separate "gene pools" for different types of strategies. For each type of strategy the system will derive the offsprings by operating on the whole pool to which they belong.

Once a company, represented by an agent, goes below a certain profit margin, it will be taken out of the market. In return the system will eventually create a new strategy out of the selection of existing strategies, weighted by the representation of the corresponding strategy in the market. The parameters of a newly created strategy instance will be chosen based on the gene pool of the corresponding strategy. This process is stochastic in nature and represents a mixture of reproduction and mutation processes.

Completely new types of strategies will be created by a human. These new types of strategies will enter the market with their own gene pools. Their "children" will then again be created based on these pools. The mutation as well as introduction of completely new strategies are crucial, since they prevent the market from stagnating.

To make sure that some presently unsuccessful strategy is given a chance to conquer the market in a more favorable time, we will maintain a "repository" of all strategies that were washed away from the market and randomly reintroduce them.

4. A TEST MODEL

To illustrate our proposed approach we designed a simulated society of customers and suppliers who live and interact in a city. The city is a circle of radius R. Customers appear in the city in intervals governed by a stationary Poisson process with a fixed frequency λ^{c}:

$$t^{c}_{i+1} = t^{c}_{i} - \frac{1}{\lambda^{c}} \log U[0,1],$$

where time is continuous and infinitely divisible. The distribution of customers is assumed to be in the equilibrium, so the society of suppliers should evolve to meet the distribution of customers. Customers appear on the market according to the following rules expressed in polar coordinates:

$$\begin{aligned} r &\sim U[0,R] \\ \alpha &\sim U[0,2\pi] \end{aligned}$$

where U is a random variable with uniform distribution. Several different types of suppliers are modeled by different sizes of their "factories". Bigger factories have lower production costs. Each factory is capable to serve one or more customers per unit of time at a cost c^{work}, or stay idle for any interval at some cost c^{idle} per unit of time. Suppliers are introduced in the market by rule similar to the one used for customers:

$$t^{s}_{i+1} = t^{s}_{i} - \frac{1}{\lambda^{s}} \log U[0,1],$$

Each supplier is granted some price level on entry and it maintains this level of price during its lifetime. Upon entry, a customer observes a selection of suppliers and chooses the one that offers the greatest benefit, where the benefit is a linear function of the supplier's price, distance to the customer, and time delay due to scheduling of other customers' tasks.

Each supplier is audited at regular time intervals and dismissed from the market if its wealth happens to fall below zero.

Price levels of the same size suppliers are considered to be a gene pool of the particular suppliers' type. We also assume that the structure of a gene pool of some type depends on the distance from the center of the city.

Every once in a while the structure of gene pools is recalculated as a function of type and distance. At the same time the density of the population is updated as a function of distance, and a new distribution of strategies by types is calculated. To smooth the effects of the limited society population, all changes enter the above described distributions with a "learning rate" γ.

4.1 Expectations

We expect the simulation to exhibit some patterns of gene pools adjustment to the market situation. It is likely that

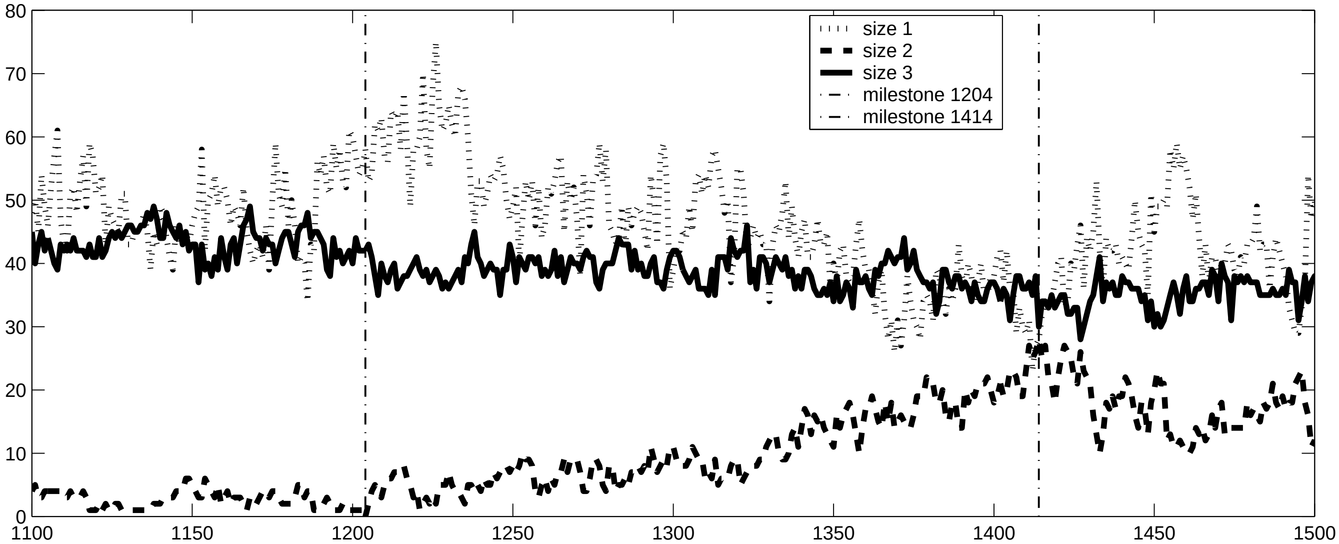

Figure 2: Population for the time period between the milestones 1100 and 1500 of city simulation.

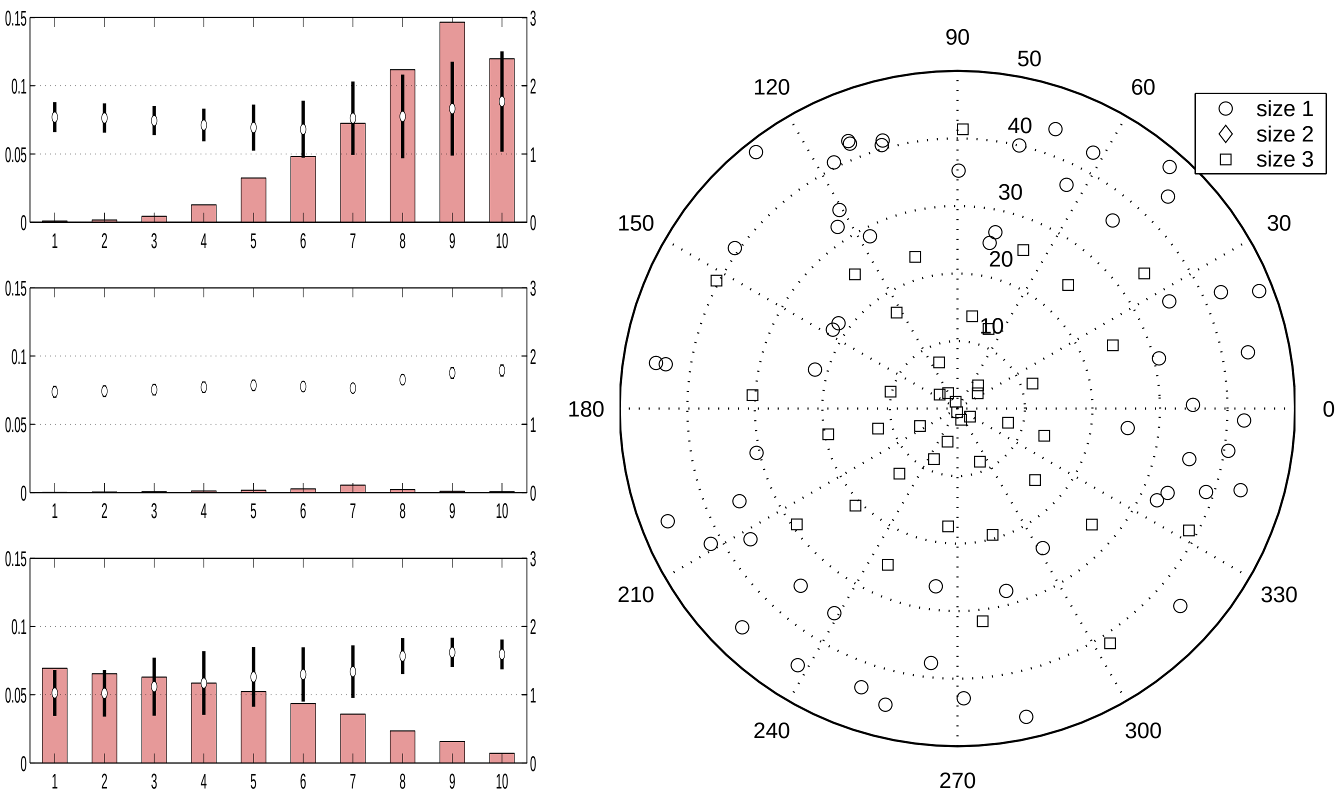

Figure 3: City simulation results after milestone 1204: gene pools (left) and city (right).

with time the relative sizes of populations of different supplier types will change, as will the patterns of average suppliers' prices and densities. We also expect that suppliers of large size should perform better near the center of the city, since their costs are lower and there are many more customers in the neighborhood. Smaller suppliers will survive better on the boundaries, where large suppliers will not have enough customers. On a final note, the higher level of competition in the center should drive the prices and the profit margins down.

4.2 Simulation Results

To verify our expectations we conducted several experiments with a variety of initial conditions. The results of one of these simulations are shown in Figures 2-4.

Figure 2 displays the population of different supplier types

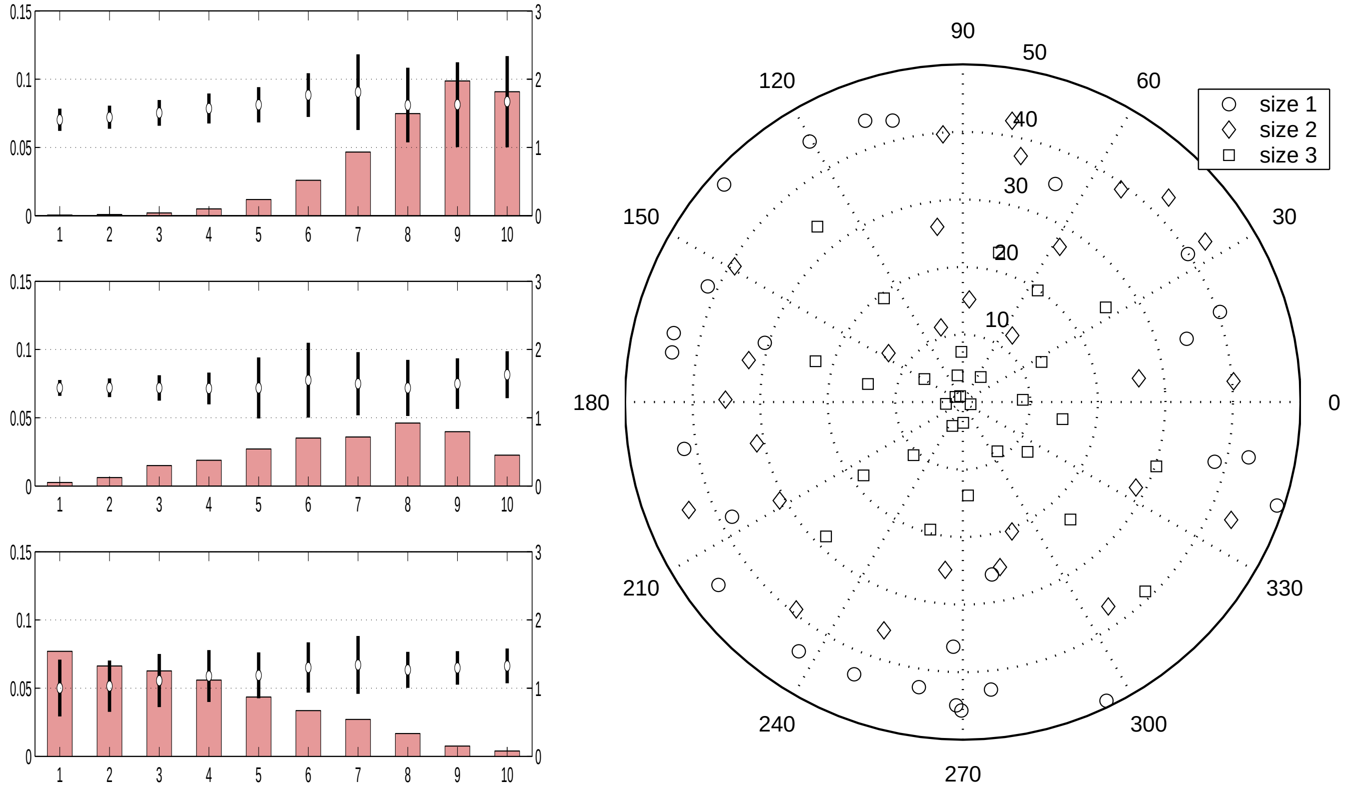

Figure 4: City simulation results after milestone 1414: gene pools (left) and city (right).

as a function of milestones. Each milestone stands for one million transactions in the market. In the figure the x-axis represents the milestones and the y-axis represents the population of each particular type. We observe that at milestone 1204 suppliers of size 2 have been entirely taken out of the market, due to their negative profit at the audit time.

On the other hand, we see that suppliers of the same size reenter the market shortly after they have been completely removed from the market. The reason is with some noise factor (set at 5% for these experiments) a strategy of a newly created supplier is chosen at random among all present and retired strategies. Hence every retired strategy has a change to enter the market again.

This models the case where a new supplier tries to enter the market with a strategy that is currently not in the market. This is an important property, because it mirrors events from the real world. Some suppliers might enter the market just for testing out some existing strategies. Alternatively, suppliers can come up with new strategies, and introduce them to the market, because they believe they have a winning strategy in the long run and they want to try it out.

We can see this happening in the timeline with the shape of the curve of supplier agents of size 2 in Figure 3. These suppliers reentered when the market situation turned favorable and increased their market share drastically to a maximum (about 25% market share) at milestone 1414. This is an important result, since the introduction of noise allows retired strategies to reenter the market in a favorable situation. The more share size 2 suppliers gained, the more were lost proportional to that by size 1. Size 3 suppliers also lost somewhat, but not as much.

Figure 3 shows the state of the city just at the demise of the strategy to own a factory of size 2 (milestone 1204). The left part of this figure shows three gene pools for factory sizes 1 (top), 2 (middle) and 3 (bottom).

In each of the gene pool graphs the x-axis shows ten concentric city zones numbered starting from the center, the left y-axis and histogram bars show the size of the population of the corresponding strategy in the particular zone relative to the whole population, and, finally, the right y-axis and error bar graph represent average values and standard deviations of profit margins.

The right part of Figure 3, in turn, gives a "bird eye" snapshot of the city at the given time point. We can see that suppliers of size 1 and 3 have divided the city into"zones of control."

It can be seen from Figure 3 that factory size 3 suppliers tend to operate near the center of the city, while size 1 suppliers prefer outer city zones. This behavior is similar to what we expected, although a picture of profit margins is not very clear. To get a better picture of the prices and profit margins we consider the state of gene pools after milestone 1414 in Figure 4.

The distributions of the population and profit margins correspond to our intuition. The comeback of suppliers which own a factory of size 2, caused the suppliers of size 1 to increase the average price in zones 4, 5, 6, 7, and 8, and at the same time suppliers of size 3 to reduce their prices in the zones 7, 8, 9, and 10. Size 2 supplier agents have found their appropriate niche in these zones over the particular time period. It is also important to note that, although the gene pools reached a relatively stable state, the population shares fluctuate all the time, as shown in Figure 2.

Using the evolutionary environment we have the advantage that at any point of the simulation we have access to fully specified strategies.

4.3 Future Work

At this stage we have mostly looked into simple pricing strategies, that conform to our knowledge of common market situations. This set of experiments was meant to show that the evolutionary framework gives repeatable and reliable results.

The next step is to create more different supplier types, where each type consists of a strategy (so far price was the only strategy parameter) and a particular size.

As an example, we could create suppliers that are discount seekers. Those enter the market in a particular location, after they have sampled the city and found the maximum price in the location around it. To do that, they would pretend to be a customer and ask for a price at a different location in the city. After they have found the maximum price, they would give a discount to this price to compete with the other suppliers in the neighborhood.

Future work will also include other aspects of agents' interactions, such as the study of effects of coalition formation. In [25] coalition formation is studied for agents that are cooperative, while MAGNET agents are self-interested.

5. COMPARISON OF DIFFERENT METHODS FOR EXPERIMENTATION

With the proposed evolutionary framework we hope to design an environment for studying effectively and efficiently properties of multi-agent systems.

An evolutionary environment is just one way to do experimentation in a multi-agent system. Another approach would be to set up a competition, where different researchers would develop their own agents and let them compete with agents written by others. An example of such an approach is the Trading Agent Competition (TAC) [1, 2].

5.1 Comparison with the Trading Agent Competition

There are similarities between competitions and the evolutionary environment we proposed. TAC provides a forum in the domain of e-marketplaces to compare approaches (agents) from a diverse collection of agent developers: *"TAC trading agents operate within a travel shopping scenario, buying and selling goods to best serve their given travel clients. TAC scores the results based on the client's preferences for trips assembled, and net expenditures in the travel markets."* [excerpt from TAC-01 description].

In the following we outline similarities and differences between the TAC environment and the evolutionary environment that we propose.

Open competitions are becoming common in many research areas. Competitions motivate researchers to participate, and incentives, like winning and publishing a paper afterwards, stimulate researchers to come up with a good piece of work. The more different research groups participate, the more diverse the population of agents becomes and the pool of available strategies grows.

Competitions are designed to provide a fair environment. Every year the rules are somewhat changed so that groups which have already competed will not have such a big advantage. Furthermore, the seeding data (clients) are generated randomly, so that all groups have a fair distribution of different types of clients.

The evolutionary environment we propose is fair as well. Suppliers which have unsuccessful strategies will eventually get removed from the market by the auditor and agents that make profit with their strategies will stay in the market.

The TAC competition provides a platform (servers, APIs, and the protocol for communication between the server and the agents) to deploy agents. This spares the difficult job of creating an environment on one's own. Researchers just need to encode their strategies (depending on the task) into their agent and plug it into the environment.

Multiple runs during the competition allow to gather a large set of training and test data. As a result of gathering lots of data, agents are also able to apply machine learning techniques in their strategies and get in this way a competitive advantage [27].

TAC characteristics:

- The competition takes place only once a year. This is a long time period to test if the performance of your algorithms has improved.
- The free disposal assumption offers space for the design of interesting and complex strategies.
- A drawback of this competition is that each group has a predefined number of rivals and clients. This limits the kind of strategies an agent can have and suggests specialization of an agent to some fixed market conditions.
- Non-transitivity between agents. During the rounds of the competition agents will loose against other agents. Some of these agents would have been able to win against their enemies in a free market in the long run, since a real environment is highly dynamic, and small periods of time are not so important.
- The competition might not be of interest to every research group. In this case the alternative is to build your own an environment for experimentation.

Evolutionary environment characteristics:

- Controllable strategies and data collection. This offers a way of doing systematic testing of agent strategies. The setup allows one to fix some strategies while others can vary. This allows drawing conclusions about certain kinds of strategies, without making them too complex. For instance, we could fix the strategy of the customer and experiment with different strategies from suppliers.
- The range of problems which can be studied is larger than in TAC. The evolutionary environment offers more opportunities to employ new strategies, such as customer and supplier agents, compared to TAC.
- Testing over a long-period of time is possible. As opposed to TAC, where the competition takes place in a short period of time, in the evolutionary environment one is able to run long term experiments over weeks or even months. This continuous setup allows to observe phenomena which can not be studied over short periods of time. For example, some strategy which was removed from the market at some point can re-enter the market at a later time and be successful. This

kind of invasion is possible, since the agent strategies are dependent on other agents which are currently in the environment, and those will change constantly.

- The type and number of customers' and suppliers' agent change frequently in an evolutionary fashion. This brings this environment closer to a real economic market and reduces the likelihood of collusion between agents.
- Reputation building is a vital part of any real system. Throughout repeated interaction the agents build their reputation based on their profit and their ability to keep their commitments. The more tasks an agent fulfills, the more is its profit and the higher is its reputation. In the proposed evolutionary environment all contacts between agents are voluntary and dynamic, hence reputation building is of major importance for agents.

5.2 Comparison with Other Methods

Research has been done in designing bidding strategies and assessing their performance. Kephart, Hanson, and Greenwald have written a survey article aimed at understanding collective interactions among agents that dynamically price services or goods [14]. They discuss and compare several pricing strategies.

Examples of price-wars caused by agents that dynamically price their information bundles are described in [13]. The data used for the experiments are not real data, but are generated synthetically making some economic assumptions and using random distributions. Because of the complexity in analyzing experimental results, experiments are limited to two agents.

Understanding collective interactions among agents that dynamically price services or goods is discussed in [14], where several pricing strategies are compared. However, no framework for experimenting with the strategies is proposed.

A simulation based approach to study dynamic pricing strategies in finite time horizon markets is described in [9]. The study is conducted using a market simulator, called the Learning Curve Simulator, as a tool for discovering the factors that determine successful market strategies. The study focuses on a finite market, i.e. a market with a finite time horizon, seller inventory, and buyer population.

The strategy used by the seller makes no assumptions about the behavior of the buyers or the type of buyers in the marketplace, it simply tries to respond to changes. Buyers stay in the market until either they have purchased a good or their lifetime has expired. Sellers can adapt their strategy every day.

A bidding strategy for continuous double auctions based on stochastic modeling is proposed in [20], with experimental results obtained by simulating the evolution of the agent population as they adapt their strategy by observing what happens in the environment.

There are various attempts to model very large multi-agent systems at the macroscopic level using physics-based methods. Shehory [26] models large scale multi-agent systems using a method based on classical mechanics. The method requires a measure of distance to the goal. Goal satisfaction is modeled by particle collisions between dynamic particles, the agents, and static particles, the goals. Most of the examples presented involve physical agents that operate in a 2D environment, where Euclidean distance is an obvious choice as the distance measure.

Similarly, Lerman [16] proposes a general methodology for mathematical analysis of multi-agent systems. The analysis is limited to systems that obey the Markov property, i.e. such that the agent's future state depends only on its present state. This is not the case in MAGNET.

5.3 Creation of Test Suites

Leyton-Brown et al [17] proposed a universal test suite for winner determination algorithms in combinatorial auctions. Their work provides well-understood test cases for comparing the performance of algorithms. The test suite currently does not include cases with precedence and time constraints and, thus, is not directly applicable to the MAGNET framework.

Test suites are important to compare performance of algorithms, but do not always capture the complexity of the domain as our proposed evolutionary system does.

We have performed systematic studies [7] to characterize the performance of the winner determination algorithms that we have developed for MAGNET. Our study follows the methodology outlined in [12].

We were interested in three measures of performance: speed, scalability, and predictability. Speed and scalability are important because combinatorial auction winner determination is known to be $\mathcal{NP}$-complete and inapproximable [23]. It scales exponentially with the number of tasks and at best polynomially with the number of bids. However, the addition of temporal constraints makes the MAGNET winner-determination problem scale exponentially in the number of bids as well. We have reported full details on our experimental results in [6].

The evolutionary framework we are proposing here complements that work by providing a richer way of evaluating different strategies, by assessing how strategies affect the long-term survival of agents, and by observing the effects of interactions of strategies in the market.

6. CONCLUSIONS

Complex system with many parameters and with stochastic properties are difficult to assess. Multi-agent marketplace systems, where agents can enter and leave the market at any time are specially hard to analyze because the agent strategies depend on the behaviors of other agents. Yet, there is no standard method for supporting systematic experiments in such systems.

We have proposed building an evolutionary system with a setup that helps the system reach a dynamically stable condition, and where agents and strategies are introduced and allowed to adapt or perish.

The outcome of using an evolutionary system is to produce several different strategies, not only an optimal one. Strategies that survive could vary all over the spectrum, from strategies that are very fast but expensive, to simple to compute strategies with long delivery delays, to strategies that depend on the size of the company, etc.

In our experiments we distributed initially all the strategies to all different sizes of enterprises, but during the simulation of evolution certain strategies drifted towards certain kinds and sizes of enterprises. The evolutionary framework allows us to observe how new behavior patterns evolve over time, and how new strategies are introduced seamlessly.

7. ACKNOWLEDGMENTS

Partial support for this research is gratefully acknowledged from the National Science Foundation under awards NSF/IIS-0084202 and NSF/EIA-9986042.

8. REFERENCES

[1] Trading agent competition 2001. http://auction2.eecs.umich.edu/, 2001.

[2] Trading agent competition 2002. http://www.sics.se/tac/, 2002.

[3] R. Axelrod. *The complexity of cooperation*. Princeton University Press, 1997.

[4] R. M. Axelrod. *The evolution of cooperation*. Basic Books, 1984.

[5] J. P. Carpenter. Evolutionary models of bargaining: Comparing agent-based computational and analytical approaches to understanding convention evolution. In *Computational Economics*, pages 25–49. Kluwer Academic Publishers, 2002.

[6] J. Collins. *Solving Combinatorial Auctions with Temporal Constraints in Economic Agents*. PhD thesis, University of Minnesota, June 2002.

[7] J. Collins, M. Gini, and B. Mobasher. Multi-agent negotiation using combinatorial auctions with precedence constraints. Technical Report 02-009, University of Minnesota, Department of Computer Science and Engineering, Minneapolis, Minnesota, February 2002.

[8] J. Collins, W. Ketter, and M. Gini. A multi-agent negotiation testbed for contracting tasks with temporal and precedence constraints. *Int'l Journal of Electronic Commerce*, 7(1):35–57, 2002.

[9] J. M. DiMicco, A. Greenwald, and P. Maes. Dynamic pricing strategies under a finite time horizon. In *Proc. of ACM Conf on Electronic Commerce (EC'01)*, October 2001.

[10] S. Forrest. Genetic algorithms: Principles of natural selection applied to computation. *Science*, 261:872–878, 1993.

[11] R. J. Gaylord and L. J. D'Andrea. *Simulating Society*. Springler-Verlag, 1998.

[12] H. H. Hoos and T. Stützle. Evaluating Las Vegas algorithms – pitfalls and remedies. In *Proc. of the 14th Conference on Uncertainty in Artificial Intelligence*, pages 238–245. Morgan Kaufmann Publishers, 1998.

[13] J. O. Kephart and A. R. Greenwald. Shopbots and pricebots. In A. Moukas, C. Sierra, and F. Ygge, editors, *Agent Mediated Electronic Commerce II*, volume LNAI1788. Springer-Verlag, 2000.

[14] J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, 2000.

[15] J. T. Landa. Bioeconomics of some nonhuman and human societies: new institutional economics approach. In *Journal of Bioeconomics*, pages 95–113. Kluwer Academic Publishers, 1999.

[16] K. Lerman. Design and mathematical analysis of agent-based systems. In *Lecture Notes in Artificial Intelligence (LNAI 1871)*, pages 222–233. Springer Verlag, 2001.

[17] K. Leyton-Brown, M. Pearson, and Y. Shoham. Towards a universal test suite for combinatorial auction algorithms. In *Proc. of ACM Conf on Electronic Commerce (EC'00)*, pages 66–76, Minneapolis, MN, October 2000.

[18] K. Lindgren. Evolutionary dynamics in game-theoretic models. In *The Economy as an Evolving Complex System II*, pages 337–367, 1997.

[19] R. R. Nelson. Recent evolutionary theorizing about economic change. *Journal of Economic Literature*, 33(1):48–90, March 1995.

[20] S. Park, E. H. Durfee, and W. P. Birmingham. An adaptive agent bidding strategy based on stochastic modeling. In *Proc. of the Third Int'l Conf. on Autonomous Agents*, 1999.

[21] C. Phillips and M. Meeker. The B2B internet report – Collaborative commerce. Morgan Stanley Dean Witter, April 2000.

[22] D. Rode. Market efficiency, decision processes, and evolutionary games. Department of Social and Decision Sciences, Carnegie Mellon University, March 1997.

[23] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135:1–54, 2002.

[24] T. C. Schelling. *Micromotives and Macrobehavior*. W. W. Norton & Company, Inc., 1978.

[25] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200, 1998.

[26] O. Shehory, S. Kraus, and O. Yadgar. Emergent cooperative goal-satisfaction in large scale automated-agent systems. *Artificial Intelligence*, 110(1), May 1999.

[27] P. Stone, R. E. Schnapire, M. L. L. J. A. Csirik, and D. McAllester. ATTac-2001: A learning, autonomous bidding agent. Submitted to the Eigtheenth National Conference on Artificial Intelligence (AAAI-2002), January 2002.

[28] L. Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. ISU Economics Working Paper No. 1, Department of Economics, Iowa State University, December 2001.

[29] J. W. Weibull. *Evolutionary Game Theory*. The MIT Press, 1995.