AN EVOLUTIONARY FRAMEWORK FOR LARGE-SCALE EXPERIMENTATION IN MULTI-AGENT SYSTEMS*

Alex Babanov, Wolfgang Ketter, and Maria Gini
Department of Computer Science and Engineering, University of Minnesota

Abstract We discuss a construction of an evolutionary framework for conducting large-scale experiments in multi-agent systems for applications in electronic marketplaces. We describe how the evolutionary framework could be used as a platform for systematic testing of agent strategies and illustrate the idea with results from a simple supply-demand model. We further explain how to integrate the proposed framework in an existing multi-agent system and demonstrate our approach in the context of MAGNET, a multi-agent system where agents bid over complex combinations of tasks with time and precedence constraints.

Keywords: Multi-agent systems, economic agents, evolutionary methods, simulation.

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 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 calculations to promptly provide a human decision maker with a refined list of alternatives. However, we believe that to exploit the true potential of electronic marketplaces, software agents need to be able to make their own decisions and adapt their strategies to the current situation.

*Work supported in part by the National Science Foundation, awards NSF/IIS-0084202 and NSF/EIA-9986042.

A major difficulty that hampers the acceptance of software agents as decision makers is the lack of systematic and accepted methods to assess and validate the agents' decisions in a multi-agent system. We are not concerned here with the broad issue of software validation. We assume that proper software design and testing methods are used in the development of the software agents. Our concern is with the methods (or lack thereof) to assess and validate the strategic decisions agents make, and their ability to adapt to changing market situations.

An issue in assessing multi-agent systems is that there is not enough real-world data available to perform comprehensive testing. At the same time, analytical modeling for the majority of less than trivial problems is prohibitively hard.

In this paper we propose to design a large-scale test environment based on an evolutionary approach to economic simulation. We specifically address the question of how to assess agent strategies in an ever changing and heterogeneous market environment.

We start by proposing in Section 2 an evolutionary approach, and we support the proposal with experimental results obtained from a simple supply-demand model. We then consider in Section 3 practical issues of building an evolutionary testing environment on top of an existing Multi-Agent System (MAS). Finally, in Sections 4 and 5 we compare our proposed approach with other existing methods and we outline future work.

2. An Evolutionary Framework for Large-scale Experimentation

A major obstacle in the way of understanding the properties 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 to make up for the scarcity of data.

The rationale behind our choice of an evolutionary framework is that it is able of revealing patterns of macroscopic behavior in a society of agents without requiring a complex theory of agent optimization criteria, or of strategic interaction [23].

The methodology we propose is essentially derived from the application of evolutionary techniques to game theoretical problems [18, 22, 26]. The evolutionary game theory studies equilibria of games played by populations of players, where players are myopically rational and have conflicting interests.

The “fitness” of the players derives from the success each player has in playing the game governed by the natural selection. Agents which do not perform well, because of their strategy, will eventually disappear from the market. In the case of electronic markets the players are customer and supplier agents, and their fitness is determined by the strategies they use to secure their profit.

2.1 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 resolution to this issue is to describe the agents’ strategies in terms of gene sequences and to use cross-breeding and mutations to ensure the desired diversity.

In large-scale multi-agent systems agents can employ a variety of methodologies, such as Q-Learning, Neural Networks, Game Theoretic models, Genetic Algorithms and others. It is hard to imagine that each and every one of the strategies that are based on the above mentioned methodologies 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 difficult to come up with an encoding for even well studied problems [13], let alone complex domains such as electronic markets.

Our proposed approach to the problem described above 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. In our test model, which is described and examined in the following sections, an information pool is derived from statistical data.

A company, represented by an agent, that receives negative profits over a certain period of time, is taken out of the market. In return the system eventually creates a new company with a variant of one of the existing strategy types. When a new company is created, the probability of selecting a particular type of strategy for that new company is weighted by how represented the strategy is in the current market. The parameters of a newly created strategy instance are chosen based on the gene pool of the corresponding strategy.

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

Completely new types of strategies can be created by a human. These new types of strategies enter the market the same way as the "retired" strategies, i.e. they are added to the list of available strategies, in the hope of acquiring a noticeable market share as soon as the market conditions become favorable.

2.2 Test Model

Our test model is a continuous time discrete-event simulation of a society of suppliers of some service and customers, which live and interact in a circular city of radius R. Customers appear in the city in intervals governed by a stationary Poisson process with a fixed frequency λ^{c}:

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

where $U[x,y]$ is distributed uniformly on the interval $[x,y]$. The distribution of customers is intentionally fixed, so that the society of suppliers had to evolve and match it. Customers appear on the market according to the following rules expressed in polar coordinates:

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

Several different types of suppliers are modeled by different sizes of their "factories." Bigger factories have increasingly lower production costs. Suppliers are introduced to the market by rule similar to the one used for customers.

A new supplier enters the market with a fixed price of its service. Every supplier is audited at regular time intervals and dismissed from the market if its profit becomes negative. These rules ensure that although each particular supplier cannot adjust its price, the society of suppliers employing the same strategy will eventually evolve to find the right price by loosing its least successful members.

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.

The probability that a supplier of a particular type will enter the market next is proportional to the number of suppliers of its type that are surviving in the market. Another possibility to enter the market is through a small noise factor (set at 5% for the experiments in this paper). With a probability equal to the noise factor a strategy of a newly created supplier is chosen at random among all present and retired strategies.

Hence every retired strategy has a chance to enter the market at a more favorable time. The noise also provides a way for completely new types to enter the market, as it will be shown later in the experimental results.

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.

2.3 Expectations

We expect the simulation to exhibit some patterns of gene pools adjustment to the market situation. It is most likely that the relative sizes of populations of different supplier types will change with time. The price distribution and the density of the suppliers as functions of the strategy type and the distance from the center are likely to adapt as well.

It is reasonable to expect that large size suppliers would perform better near the densely populated and therefore highly competitive center of the city, because of their lower production costs. Smaller suppliers will survive better on the boundaries, where transportation costs become increasingly important to compare with the cost advantage of the large suppliers. Consequently, the higher level of competition should drive the prices and the profit margins down close to the center.

The reasonably evolving society of suppliers should adapt to the dynamic changes in the parameters of the customer distribution. Consistently with the expectations outlined before, the increased frequency of the customers' arrival should increase the "habitat" of the large suppliers, while the opposite change should give a leverage to the small ones.

On a final note, new and retired supplier types should be able to acquire a position in an existing market by the means of noise in the supplier type selection process.

To verify our expectations we conducted several experiments with a variety of initial conditions. Two representative experiments are considered in the following.

2.4 Simulation Results: Noise Factor

In the first experiment the simulation starts out with suppliers of size 1 and 2. After some time suppliers of size 3 enter the market through the 5% noise factor, meaning that initially every new supplier has a chance of about 1.67% to enter the market with size 3 factory. This experiment models a situation when a new strategy is designed and some suppliers try to enter the market with its benefit (or, alternatively, some of the existing supplier decide to switch).

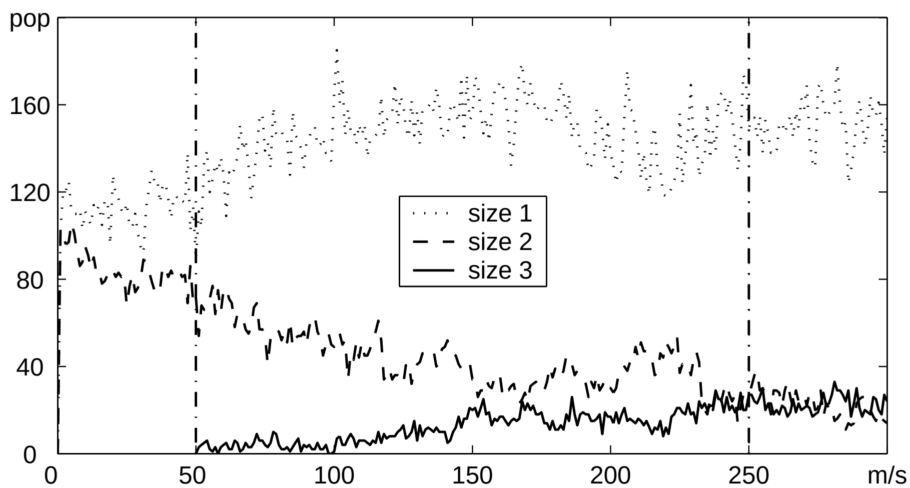

Figure 1. City population for different supplier types as a function of milestones. The vertical lines denote milestones 50 and 250.

Figure 1 displays the population of different supplier types as a function of milestones. Each *milestone* (m/s) stands for two 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. Suppliers of size 3 enter the market at milestone 50 and struggle to find their place.

After some time in the market, the size 3 supplier type proves itself to be competitive with size 2 suppliers. The more market share size 2 suppliers gained, the more were lost proportional to that by size 2, until a dynamic equilibrium with approximately same populations of both sizes was reached around milestone 250.

To reveal the mechanics of size 3 successful entry we examine the state of the city at milestones 50 and 250.

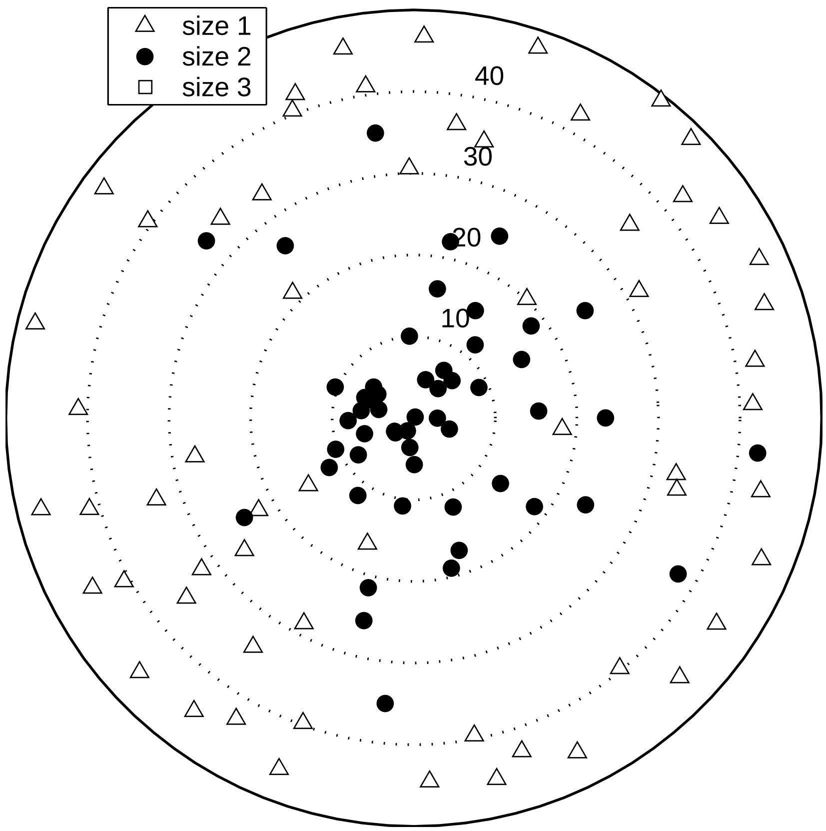

Figure 2. Distribution of the different supplier types in the city at milestone 50.

Figure 2 shows the state of the city just before the introduction of a strategy to own a factory of size 3 (milestone 50). There are two important observations to be made from this figure. Firstly, size 1 suppliers dominate the market and, indeed, as we expected tend to the rim, while the larger size 2 suppliers operate mostly in the middle of the city. Secondly, the distribution of suppliers is quite uneven with dense clusters and wide open areas situated at the same distance from the center.

Figure 3, in turn, gives a “bird eye” snapshot of the city at milestone 250. We can see that suppliers of size 3 pushed size 2 suppliers out of the center, while suppliers of size 1 still successfully inhabit the outer city zones.

The left part of the Figure 4 shows the state of the two gene pools for factory sizes 1 (top) and 2 (bottom) at milestone 50. The right part of this figure shows the gene pools for all three sizes, starting with size 1 at the top to size 3 at the bottom at milestone 250.

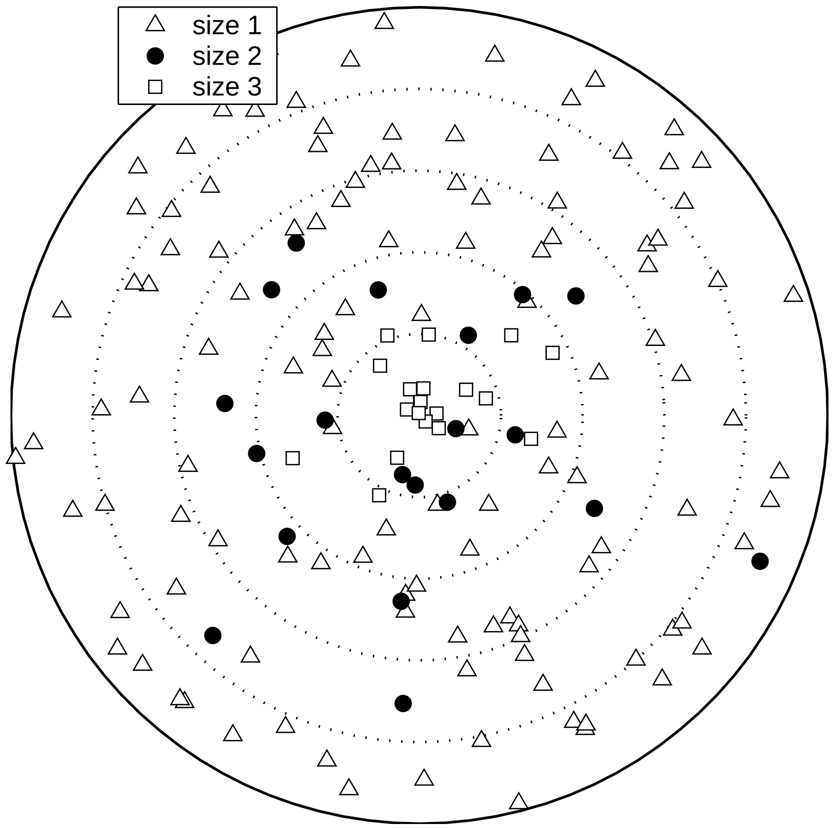

Figure 3. Distribution of the different supplier types in the city at milestone 250.

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.

It can be seen from Figure 4 that size 2 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 at milestone 250 in the right part of Figure 4.

The introduction of size 3 suppliers caused the suppliers of size 1 and 2 to decrease the average price in all zones. Size 3 supplier agents have found their appropriate niche in zones one to three. We observe that size

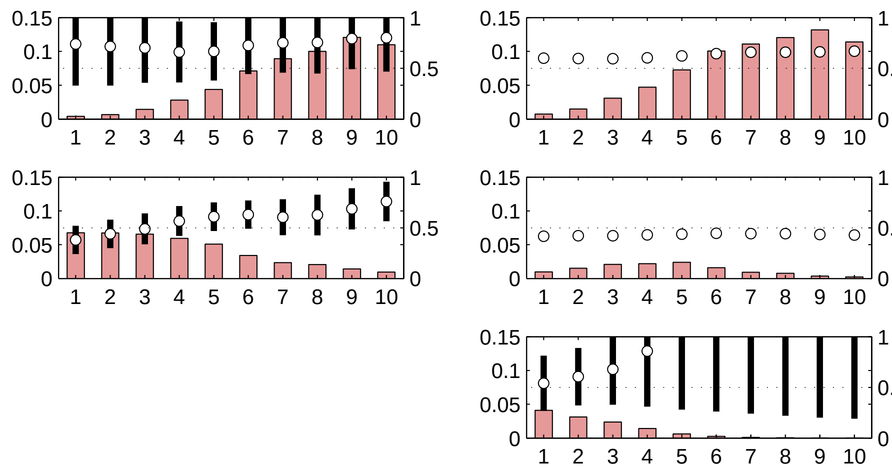

Figure 4. Structure of the gene pools of the city simulation as a function of city zones at milestone 50 (left) and at milestone 250 (right). Supplier size: 1 (top), 2 (middle) and 3 (bottom).

1 and 2 suppliers converged to stable averages prices for their services in all zones, as the variance is very low.

It is also important to note that, although the gene pools reached a relatively stable state, the population shares continuously fluctuate, as shown in Figure 1. Also, the high variance of size 3 profit margin implies that the market state may change in a future as this strategy try to find the right price distribution.

2.5 Simulation Results: Changing Environment

In the second experiment we reduced the frequency of customers' arrival by 1/3 of its initial value half way through the simulation. This was meant to emulate the loss of interest in the supplied service due some economic factor, such as depression or introduction of the alternative service.

One of the results of the experiment is depicted in Figure 5. The figure shows the entry probabilities for each of the three supplier types. In this figure we can observe two important effects. First and foremost we see that the market reaches a relatively stable state shortly both after the beginning of the simulation and after the change of conditions. Secondly, we observe that after the change the size 3 suppliers loose a sizable part of their market share to size 1 suppliers. Lower frequency

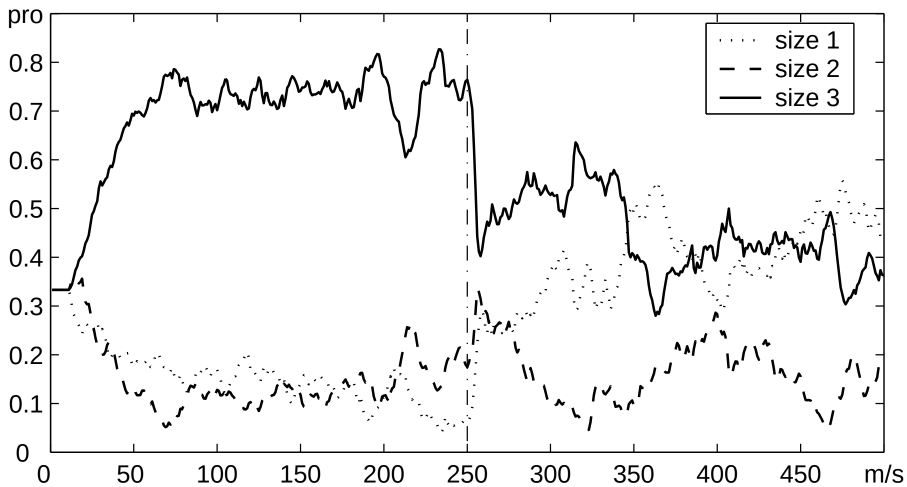

Figure 5. Probabilities of a new supplier entry for different types as a function of milestones. The customer frequency was reduced by a factor of 1/3 from its initial value at milestone 250.

of customer arrival resulted in the disadvantage for large suppliers in accordance with our expectations.

3. Integration of the Evolutionary Approach into an Existing MAS

In this section we demonstrate a way to add an evolutionary framework to an existing MAS. For the example purposes we use MAGNET, a multi-agent system we have designed to study agents in auctions for combinations tasks with time and precedence constraints [9]. In essence, MAGNET is a mixed-initiative system, in which intelligent software agents facilitate the deliberation process of human decision makers. It is possible, however, for research purposes, to exclude a human from the loop and let the agents select autonomously the best course of actions.

3.1 MAGNET Architecture

In MAGNET we distinguish between two trading agent roles, *customer* and *supplier* (see Figure 6). The customer has a set of tasks to be performed and needs to solicit resources from suppliers by issuing *Request for Quotes* (RFQs) through an agent-mediated market.

MAGNET agents participate in a first-price sealed-bid reversed combinatorial auction over combinations of tasks with precedence relations

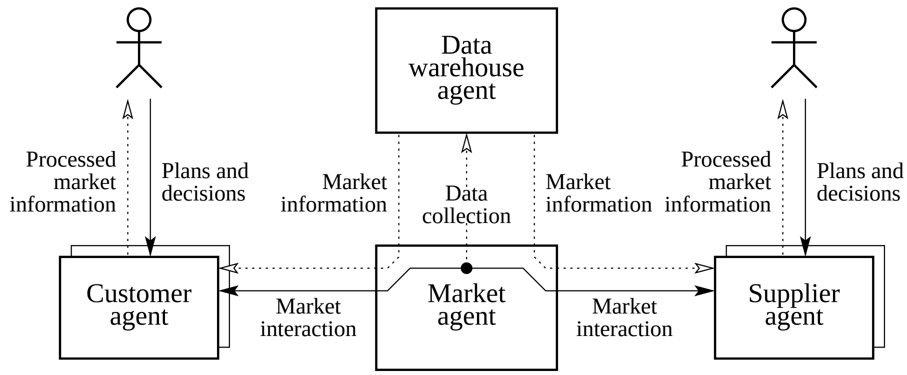

Figure 6. The mixed-initiative MAGNET architecture.

and temporal constraints. After the auction ends, the customer agent solves the winner determination problem and awards bids.

The market agent is responsible for coordination of tasks, i.e. distributing customers' RFQs among an appropriate selection of suppliers, collecting and timing bids, monitoring interactions during task execution phase, etc. The data warehouse agent collects information on the transactions, and makes it available in form of statistical data to agents and their owners.

3.2 A Practical Example

The following example of a house construction illustrates how MAGNET handles problems in its domain. Figure 7 (left) shows the tasks needed to complete the construction. The tasks are represented in a *task network*, where links indicate precedence constraints.

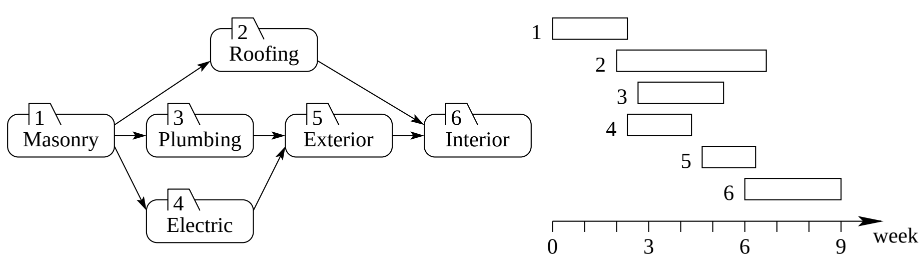

Figure 7. A task network example (left) and the corresponding RFQ (right).

The first decision the customer agent is faced with is how to sequence the tasks in the RFQ and how much time to allocate to each of them. For instance, the agent could reduce the number of parallel tasks, allocate more time to tasks with higher variability in duration or to tasks that are in short supply in the market. Presently MAGNET uses a simple CPM-based algorithm for generating RFQs [8]. An alternative approach based on the Expected Utility Theory is being actively researched by our group [3, 4].

A sample RFQ is shown in Figure 7 (right). Note that the time windows in the RFQs do not need to satisfy the precedence constraints; the only requirement is that the accepted set of bids satisfies them.

3.3 Introduction of Evolutionary Components

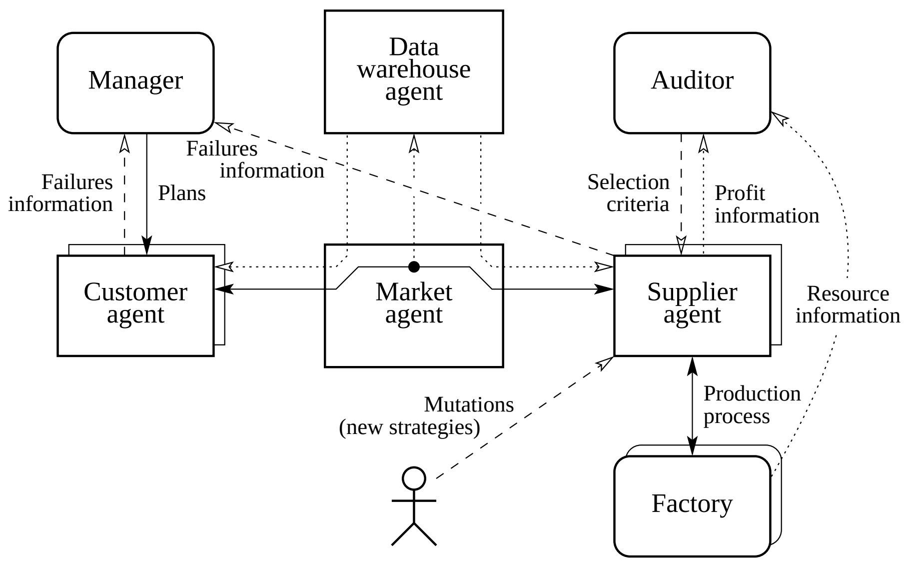

Figure 8. The MAGNET architecture adjusted to the evolutionary paradigm. The rounded boxes show specific evolutionary components.

In order for MAGNET to operate in the evolutionary framework, we need to add components that will manage evolutionary aspects of the system and exclude human decision makers from the loop. Figure 8 shows the resulting architecture and the following list summarizes the required changes:

The **Manager** generates and distributes tasks to customer agents. It observes the rate of customers' and suppliers' failures to complete their evaluations of RFQs and bids during specified deliberation periods.The manager adjusts the frequency of issuing RFQs 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 frequency of generating the RFQs determines the size of the market population.

The **Auditor** evaluates the performance of supplier agents' strategies based on suppliers' 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" strategiesand eventually tries to put them back in the market.

The **Customer agent** makes all market decisions without help from its human supervisor and reports the rate of failures to the manager.

The **Supplier agent** also operates without human supervision and reports on its computational failures to the manager. The supplier agent coordinates its resource commitments with its own factory.

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.

Human participants submit new strategies to the pool of possible **mutations**. "Mutant" strategies are introduced to the market after it reaches its dynamic steady state, i.e. after the rate of issuing RFQs by the manager eventually stabilizes.

The **Data warehouse agent** collects data the same way as in the mixed-initiative configuration. In addition, it replies to data queries from the manager, the auditor and, possibly, human observers whenever it is required.

The choice of this particular architecture is determined by the need to stabilize the size of the market and by our specific interest in supply-side strategies. To satisfy these requirements we fix the population of customer agents and let the supplier agents evolve to meet the demand. The demand, in turn, is limited from above by the computational capacity of a system that runs agents' software. In case the load is overly high, the rate of agents' failures to complete calculations will signal the manager to decrease the rate of issuing RFQs, thus effectively shrinking the market.

Whether we decide to study the demand-side strategies, the architecture might be changed to make the auditor manage the evolving population of customer agents, while the manager governs the resource availability on the supply-side. The exact choice and composition of the evolutionary components is not formalized at present, however we plan to improve the methodology as we study other prospective domains. Some of such candidate domains are considered in the following section along with the approaches that are currently used to study them.

4. Experimentation in an Evolutionary Environment

The proposed evolutionary approach to large-scale simulations is not the only possibility to makeup for the lack of readily accessible real-world data. Two other methods that are widely used in the research community are analytical modeling and competitions of software agents.

The major drawback of analytical modeling lies in its failure to embrace the complexity of the real world to any significant degree. It is often useful to employ analytical methods to study some specific and usually global properties of the highly simplified models, however, they are not very helpful when the domain of interest involves many different types of agents and possible agents' strategies.

A comprehensive review of successes and pitfalls of competitions based on experiences from the Trading Agent Competition (TAC) and RoboCup can be found in [24]. In short, the competitions of intelligent software agents proved to be a dynamic and valuable source of research material. The problematic properties include overly restrictive rules, domain-specific solutions (i.e. strategies that exploit peculiar properties of the simulation environment or wordings of the rules) and invalid evaluation criteria.

In the following we focus on the important properties of the evolutionary methodology:

- A heterogeneous multi-agent system is commonly 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 by far impossible. Using an evolutionary approach allows us to search the space of parameters in an efficient way.
- A thorough study of agent strategies requires information about the behavior of other agents in the system. The evolutionary ap-

proach solves this problem by enclosing all the agents in a self-sufficient system, where they can observe each other's behavior.

- The evolutionary approach allows for formation of complex spatio-temporal patterns of behavior on the level of groups of agents. Examples range from the emergence of cooperation in an otherwise selfish society [1, 2] with possible formation of spatial patterns of strategic interaction [17] to natural phenomena, like fish schools [15].
- Finally, a simulated evolutionary environment, like most simulations, offers facilities for controllable experimentation and systematic data collection.

4.1 The Role of Evolutionary Experimentation in Application Science

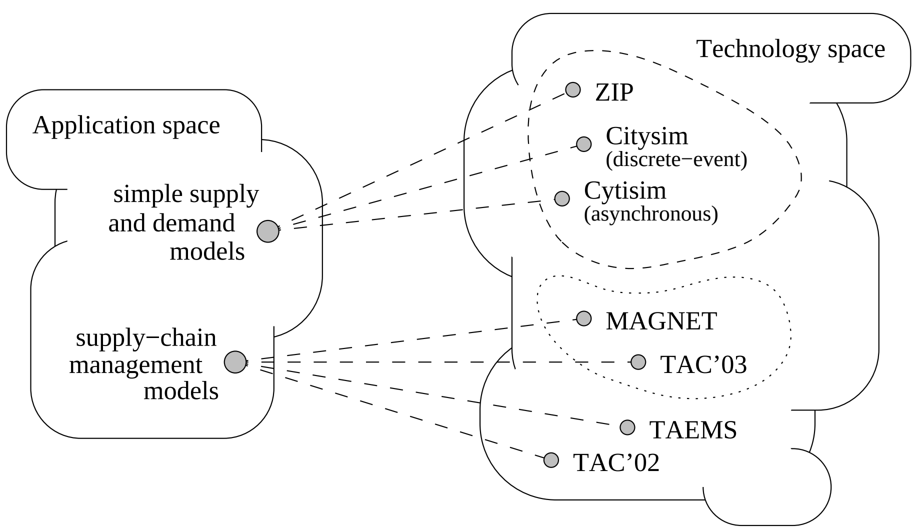

Figure 9. The division of the technology space by the degree of compatibility with the evolutionary approach.

We divide the technology space in three parts by degree of compatibility of each known or future technology to our evolutionary approach. The first subspace includes technologies that are based on one or another evolutionary methodology, such as Genetic Algorithms or Celluar Automata. Examples of such technologies include ZIP [6], our pilot discrete-event Citysim model or, perhaps, its possible asynchronous implementation.

The second collection of technologies can be restructured to be compatible with the evolutionary paradigm. MAGNET [9], as it was demonstrated in the previous section, is one perspective member of this subspace, and the new 2003 revision of TAC [25] is another.

In the third part of the space we count those technologies that are not easily convertable to the evolutionary setup, such as the older versions of TAC. Also in this subspace are technologies whose compatibility was not yet examined, e.g., TAEMS framework [11, 16].

5. Related Work

Much research has been done in the last few years in designing pricing strategies for agents, assessing their performance, and adapting to changing situations [5]. Understanding collective interactions among agents that dynamically price services or goods is discussed in [14], where several pricing strategies are compared. Examples of price-wars caused by agents that dynamically change their posted price for information bundles are described. Because of the complexity of the problem, experiments are limited to a small number of agents.

A simulation based approach to study dynamic pricing strategies in finite time horizon markets is described in [12]. The study uses a market simulator and simple strategies. The results are evaluated in terms of overall profit, but there are so many variables in the simulation that it is hard to assess the generality of the results obtained.

Continuous double auctions have been the subject of multiple studies. Cliff's [6] Zero-Intelligence Plus trader agents have minimal intelligence, yet they have been successfully used in continuous double auctions, where they performed very well even when compared to human traders [10].

The use of evolutionary methods for continuous double auctions is proposed in [19], who simulates the evolution of the agent population as they adapt their strategy by observing what happens in the environment. Cliff [7] uses genetic algorithms to learn the parameters that control how his trader agents evolve their pricing strategies. Along similar lines, an evolutionary system based on Genetic Programming is presented in [20].

The major difference with these and the work presented here, is that we are interested in understanding how strategies of individual agents interact in the market, as opposed to study specific types of auctions to learn auction rules. We are also interested in providing a methodology for studying effectively multi-agent systems with a large number of agents.

6. Conclusions and Future Work

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 behavior patterns 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.In evolutionary systems there is no fitness function, instead there is a rule which governs survival of society members based on their success. In our case, when an agent fails to make any profit for a period of time, the agent will leave the market to be eventually replaced with a more fit entity.

The outcome of using an evolutionary system around a MAS could produce several different strategies, not only an optimal one. Strategies that survive could vary from some strategies that are very fast but expensive for the customer, to inexpensive strategies with long delivery delays, to strategies that depend on the size of the company, etc. The design of the framework allows for new behavior patterns to evolve over time and for new strategy types to be introduced seamlessly.

Our future plans include examination of conditions that allow for the introduction of the evolutionary framework in other MAS and formalization of the integration procedures. We are considering applying the evolutionary framework to the year 2003 revision of the Trading Agent Competition [25], to experiment with strategies for manufacturer agents under different market conditions. We will also use the proposed approach to study strategies and patterns of strategy interaction in the context of the MAGNET system.

References

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

[2] Robert Axelrod. *The complexity of cooperation.* Princeton University Press, 1997.

[3] Alexander Babanov, John Collins, and Maria Gini. Risk and expectations in a-priori time allocation in multi-agent contracting. In *Proc. of the First Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, volume 1, pages 53–60, Bologna, Italy, July 2002.

[4] Alexander Babanov, John Collins, and Maria Gini. Scheduling tasks with precedence constraints to solicit desirable bid combinations. In *Proc. of the Second Int'l Conf. on Autonomous Agents and Multi-Agent Systems*, Melbourne, Australia, July 2003.

[5] Christopher H. Brooks, Robert Gazzale, Rajarshi Das, Jeffrey O. Kephart, Jeffrey K. MacKie-Mason, and Edmund H. Durfee. Model selection in an information economy: Choosing what to learn. *Computational Intelligence*, April 2002.

[6] D. Cliff and J. Bruten. Minimalintelligence agents for bargaining behaviors in marketbased environments. Technical Report HPL-97-91, Hewlett Packard Labs, 1997.

[7] Dave Cliff. Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double auction markets. Technical Report HPL-2001-99, Hewlett Packard Labs, 2001.

[8] John Collins, Corey Bilot, Maria Gini, and Bamshad Mobasher. Decision processes in agent-based automated contracting. *IEEE Internet Computing*, pages 61–72, March 2001.

[9] John Collins, Wolfgang Ketter, and Maria 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.

[10] Rajarshi Das, James E. Hanson, Jeffrey O. Kephart, and Gerald Tesauro. Agent-human interactions in the continuous double auction. In *Proc. of the 17th Joint Conf. on Artificial Intelligence*, Seattle, WA, USA, August 2001.

[11] Keith Decker. Taems: A framework for environment centered analysis and design of coordination mechanisms. In *Foundations of Distributed Artificial Intelligence*, pages 429–448. John Wiley & Sons, Inc., January 1996.

[12] Joan Morris DiMicco, Amy Greenwald, and Pattie Maes. Dynamic pricing strategies under a finite time horizon. In *Proc. of ACM Conf on Electronic Commerce (EC'01)*, October 2001.

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

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

[15] Janet 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] Victor Lesser, Bryan Horling, Frank Klassner, Anita Raja, Thomas Wagner, and Shelley XQ. Zhang. BIG: A resource-bounded information gathering and decision support agent. *Artificial Intelligence*, 118(1–2):197–244, May 2000.

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

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

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

[20] Steve Phelps, Peter McBurney, Simon Parsons, and Elizabeth Sklar. Co-evolutionary mechanism design: a preliminary report. In *Workshop on Agent Mediated Electronic Commerce*, Bologna, Italy, July 2002.

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

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

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

[24] Peter Stone. Multiagent competitions and research: Lessons from RoboCup and TAC. In *The 6th RoboCup International Symposium*, 2002.

[25] TAC-02. Trading agent competition 2002. http://www.sics.se/tac/, 2002.

[26] Leigh Tesfatsion. Agent-based computational economics: Growing economies from the bottom up. *Artificial Life*, 8(1):55–82, 2002.