at The Fuqua School of Business, Duke University. He
pean Institute for Advanced Studies in Management) in
and particularly using computational approaches. He has
ИBA. He is active on a number of editorial boards and has
nd management science, and seven books.

ınagement at the Kenan-Flagler Business School, Univer-
ts center on managing innovation and change, particularly
anagement and location of R&D, and managing control-
ersity of Texas at Austin. She has published in *Strategic*
f Management Journal, and *Journal of Accounting and*

Computational & Mathematical Organization Theory, 8, 221–234, 2002
© 2002 Kluwer Academic Publishers. Manufactured in The Netherlands

Emergent Specializations in a Commodity Market: A Multi-Agent Model—Graduate Student Best Paper Award, CASOS 2001 Conference

MAKSIM TSVETOVAT AND KATHLEEN M. CARLEY
Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA
email: maksim@cs.cmu.edu
email: kcarley@ece.cmu.edu

Abstract

It has been observed in many instances that markets have a tendency to segment themselves into distinct sub-markets. This paper presents a multi-agent model that illustrates emergent market segmentation. The model illustrates the way local optimization processes result in an emergent global improvement of social welfare.

Keywords: multi-agent systems, social networks, markets, computational modeling, emergent behavior, self-organizing systems

1. Introduction

With emergence of electronic transactions as a predominant form of business-to-business trading (FreeMarkets.com, 2000; FastParts.com, 2000; Chemdex.com, 1999), the study of such markets becomes more and more important. While in many ways electronic markets resemble the traditional trading floor-based marketplaces, the immediacy of available data on the behavior of such markets as a whole as well as individual market agents makes electronic markets an alluring subject of study.

It has been noted that unregulated markets have a tendency to become increasingly segmented over time (Hannan and Freeman, 1977).

A market is often described in terms of competition and natural selection. Organizations or individual may fail to flourish in certain environmental circumstances because others compete with them for essential resources. As long as the resources which sustain the market are finite and market participants have unlimited capacity to expand their business, competition must ensue.

A. Howley (1950) shows in his model that competition processes typically involve four components:

- demand for resources exceeds supply,
- selection eliminates weakest competitors
- competitors differentiate territorially or functionally, yielding a division of labor in a number of market niches,

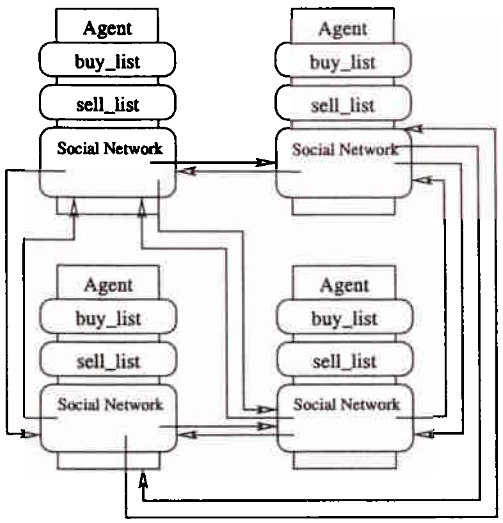

Figure 2. A set of agents.

my inventory") the agents are only concerned with their own immediate profit. Agents have no way to estimate other agents' profits or the global welfare of the market, or to predict the direction that the market will take in the future. The main goal of the agent is to execute the buy and sell orders it receives from its customers. Thus, such an agent is a fairly accurate representation of a market trader that specializes in negotiating and executing transactions but allows his or her customers to make their own buying decisions. The agent utility from each transaction is:

$$U = \mathit{TransactionPrice} - \mathit{ReservePrice} - \mathit{TransactionCost},$$

where the transaction price is the final price at the end of the negotiations and the reserve price value has been supplied by the customer.

The transaction cost depends on the specific market and can be calculated in a number of different ways:

- *TransactionCost* $=$ *const*: In this extreme case, the agents would seem to have no incentive to re-organize their communication networks, and will engage in as many communicative acts as it takes to maximize their utility. However, this case is unrealistic if we presume that the agents are subject to *bounded rationality*. In particular, processing one communicative act takes an agent a certain non-zero amount of time—and while the agents might be able to process a large number of messages, this number would still be finite during a finite clearing period. Thus, agents in the no-cost environment will still have an incentive to keep the number of messages low—with the lower bound set by a ratio of time it takes to process a message and length of the clearing period.

- *TransactionCost* $= f(m)$ *where* m *is the number of messages*: In this case, the agents have a clear incentive to optimize their communication patterns with the goal of bringing the number of messages required to complete a transaction to a minimum. They can do that by creating a social network of other agents. Stronger ties in this network are created to agents that are most likely to possess the goods in question and sell them at a fair price.

3.2. *Agent Decision-Making*

When an order comes in, a trader agent executes the transaction if the good is on *buy_list* or *sell_list* and if the reservation prices supplied by the customer are acceptable in the current market conditions.

If an agent fails to complete the transaction within the clearing period, it does not receive the positive utility, but still has to pay the communication costs. Thus, it is possible that an agent completes some of the clearing periods with a negative utility.

The agents in the market have to make a set of decisions to complete the transaction. The fitness of these decisions to the market situation is largely responsible for whether an agent will be successful (receive high utility) or not.

In each clearing period, the agents must make the following decisions:

- Which of the market goods should be traded?
- Which agents should I contact and in which order?
- Is the offer I received good enough?

The market goods are organized in a priority queue, sorted by the running-average utility derived from trading a particular good. At the beginning of each clearing period, the agent chooses an item from the list using the Metropolis criteria (i.e. an exponentially distributed random variable with λ proportional to the highest average utility among the goods). The λ parameter is functionally equivalent to the temperature of a simulated annealing search and represents the degree of entropy within each of the agents.

The result of using the Metropolis criteria is that while the agent receives approximately equal average utilities from trading different goods, the chance of one of the goods is approximately equal. However, if one of the goods is more profitable than others, the chance of it being chosen increases dramatically. As the market prices vex and wane, the value of λ changes, and probabilities of other goods being picked increase.

The decision on which agent one should talk to is done in a similar manner. Each trader maintains a table where it stores a set of values U_g, a, a running-average utility of trading good g with agent a. The utilities for each of the goods are stored in a priority queue.

The negotiation proceeds as follows:

- First, an agent chooses a good to trade.
- Using the Metropolis criteria, it then chooses a potential partner and sends a request for bids.
- On receipt of the bid, the agent evaluates the price and chooses to accept or reject the bid.
- If the bid is rejected, the agent returns to step 2 and chooses another agent to talk to.

- If a bid is accepted, both agents report the transaction to the market clearing agent and record their respective utilities, thus updating the selection tables

4. Simulations Using Multi-Agent Networks

The Multi-Agent Network Model paradigm is based upon the following assumptions:

- The simulation consists of agents
- Agents are independent, autonomous entities endowed with some intelligence
- Agents are cognitively limited
- Agents can learn knowledge about the world and referential knowledge about other agents, with a limited learning capacity
- Agents can forget
- Agents do not have accurate information about the world
- Agents do not have accurate information about other agents
- Agents communicate asynchronously and deal with complications resulting from asynchronicity (i.e. deadlocks, delays, etc) in an autonomous manner.
- Unless required by the simulation domain, there is no central mediating entity to resolve the conflicts.
- Unless required by the simulation domain, the agents do not use predefined geometrical locations or neighborhoods.

The simulation paradigm is task-independent. The task is merely defined as a function that maps a problem vector and agent's knowledge vector unto a result vector. Thus, the simulation can be easily adapted to different simulation domains, from military simulation to electronic marketplace simulation.

The simulation consists of the following parts:

Communication Infrastructure: Handles the information flow between agents, provides message buffering and manages concurrent agent processes. Provides an abstraction layer between the agent processes and physical communication—allowing a variety of options for implementing the physical communication as well as seamless clustering across several systems.

Agents: Independent processes (or threads, depending on implementation) that use the communication infrastructure to communicate with each other.

Instrumentation API: Allows researchers to collect, filter and output data about individual agents as well as the system as a whole. The instrumentation API is built upon the producer-filter-consumer model, allowing for a flexible data acquisition and output in variety of formats.

Multi-agent Network Simulations provide for a number of properties that are extremely useful in simulations of complex social systems. Among these propertiesare: (a) structural realism, (b) temporal realism, (c) information flow realism and (d) task realism.

4.1. *Structural Realism*

Artificial-life based simulations are built upon the concept of agents (or cellular automata) located on a grid of a specified shape (1-dimensional, square, toroidal, 3-dimensional, etc). Interactions are based upon the concept of proximity, defined by the agent neighborhood on the grid. Unfortunately, the choice of grid shape and type of neighborhood is arbitrary and often does not carry any face validity. Yet it can significantly alter the behaviour of the system and thus, simulation results.

While physical proximity in the real world carries some importance, the majority of human interactions are based upon social proximity—arbitrary graph structures not constrained by the concepts of grids or grid neighborhoods. To increase the face validity of social simulations, the simulation architectures must use similar structures to describe the social structure of the simulation domain.

Moreover, while the social network of an organization can be studied objectively by an outsider, none of the participants of the network actually have an accurate view of the interaction structure of the organization. They do, however, have beliefs about that structure. and use them to guide them through the interactions.

This point is important because these beliefs are often inaccurate, and change rapidly as information is processed.

In a multi-agent network simulation paradigm, the agents' interactions are governed by the formal structure of the organization, and agents' beliefs about the informal structure.

The formal structure of the organization is specified as a directed weighted graph that specifies the communication channels that are open as well as their throughput or cost of communication. The directed nature of the graph allows one to specify one-way relationships and chain-of-command relationships.

The beliefs about the informal structure are individual to every agent, and also consist of a weighted directed graph. However, when an agent joins a network, its informal relationship graph is empty, and it must learn about the informal network before it can be used for communication.

In case of the market simulation, no pre-defined formal structure exists and agents are left on their own to decide who to communicate with.

4.2. *Temporal Realism*

Most simulations are synchronous—based on the concept of time periods and turn-taking. While this provides an adequate approximation of simple interactions and games, it is not sufficient for simulation of protracted interactions and interactions where cognitive ability and cognitive load of agents are relevant to the simulation domain.

In an asynchronous system the notion of time is independent from the notion of task or communication, thus allowing realistic modelling of cognitive load and cognitive abilities of agents.

However, asynchrony presents a number of complication—including possibility of communication deadlocks and overloaded communication channels. However, such things happen in the human societies as well and thus add to the realism of the model.

4.3. *Information Flow Realism*

In a multi-agent network, the agents do not have perfect knowledge about the world. The only way to obtain information about either the world or other agents is to ask about it and then learn the results of the query—or obtain the information as a result of information exchange interaction. Agents may or may not communicate their beliefs truthfully and can be strategic about their communication. This is required for simulation of domains where the agents are competitive (i.e. electronic markets) or hostile (i.e. military simulations).

4.4. *Task Realism*

In building a multi-agent model of a social phenomenon, one has to take into a count several major factors.

First of all, the multi-agent model must accurately represent the processes present in the subject of study. In many cases it requires building a large knowledge base that mimics the cognitive processes of a human involved in the similar situation. An example of such cognitively accurate model is Soar (Jones et al., 1995), which employs complex rule-based reasoning and learning processes to emulate performance and cognitive processes of a human operator.

However, in order to create accurate models of emergent phenomena, such as market activities, one must create simulation environments that include large numbers of agents. Due to their complexity and computational requirements, Soar and other cognitively accurate models cannot be used as part of large multi-agent system. Thus, emergent behaviors are often modelled using large collections of very simple agents (cellular automata). However, the simpler agents often cannot replicate the processes involved in market environments.

Thus, in order to build scalable and yet accurate simulation of a market, one must find a balance between cognitive complexity and computational requirements of a large number of agents.

We find that such a balance can be achieved with accurately modelled transaction protocols and realistic decision functions based on agent's self-interest.

5. MarketSim—The Multi-Agent System

The MarketSim system was built on the foundation of the RETSINA multi-agent system framework. The RETSINA framework provides communication functionality (Sycara et al., 1996), yellow and white-pages services (Decker et al., 1996; Sycara et al., 1999a) that allow agents to find each other and interoperability (Sycara et al., 1999b) services. The system allows wide re-use of tools and agents across different projects.

In building this model, we have used the communication and advertising components of the RETSINA architecture, a AgentFactory (Economou et al., 2001) protocol compiler and a logging and display services agent.

6. Analysis Tools

The RETSINA system was built with the presumption that behavior of multi-agent systems as a whole can be studied in conjunction with building of the agents on a micro level. Thus, the architecture provides for an easy way to integrate global logging and visualization into the multi-agent system.

The resulting logs are usually used to visualize the agent behavior via a tool called DemoDisplay. The DemoDisplay shows each agent as an icon on a field, and shows the communications between agents using animated arrows. When two agents communicate, their icons move closer to each other, and when they stop communicating the agents retreat to their initial positions around the border of the screen.

Thus, the screen becomes a real-time display of the social network of agents, showing clusters of agents as they communicate or conduct transactions.

The output of the logging tools can be also converted into MatLab data for further analysis. In this project, the logging data is converted into agent adjacency matrices and network diagrams. It is possible to sample the network diagrams at different times within the same system, so one can analyze the dynamics of relationships between agents.

7. Virtual Experiments with MarketSim

In the beginning of the paper we have stated a hypothesis that the despite being myopic and self-interested, the agents, in their quest for higher profits, will produce a market that is optimized on global scale. With the development of the MarketSim model, it is now high time to prove the validity of the hypothesis through a set of virtual experiments.

In these experiments, we have used several markets containing 25, 50 and 100 agents. All agents are similar in the way they compute their utilities, with the exception of reservation prices at every transaction. In all experiments, transactions are done via first-price English auction.

Agents are cognitively and conversationally limited—i.e. an agent can only participate in a small number of auctions at the same time due to both communication and reasoning load. There are 3 distinct goods to be traded. An agent starts out randomly initialized to be interested in 1 or more for both buying and selling.

The timekeeping for all agents is provided by a controlled clock, allowing the experimenter to slow down the agents to an observable speed, or increase the performance to near-realtime. In runs involving larger numbers of agents, it is often necessary to slow down the simulation clock to prevent overloading of the network interfaces.

The independent variables are:

- market size—the number of agents in the market
- market saturation—proportion of agents offering each of the goods; market saturation is equivalent to social network term "initial probabilistic density".

The main measurements in each of the experiments are:

- Average Network Load—the number of messages that pass through the network in one time period

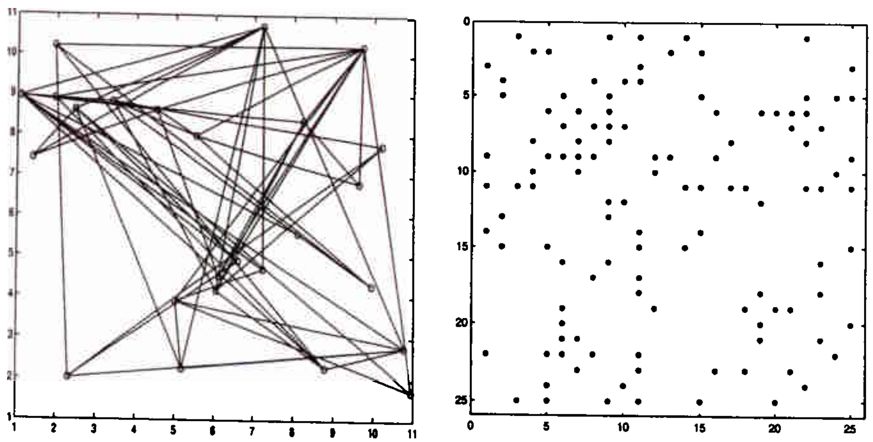

Figure 3. Social network for baseline market and its adjacency matrix.

- Average Transaction Cost—proportional to number of messages that have to be exchanged to complete one transaction
- Average Agent Utility—how well do the individual agents do?
- Overall Social Welfare—how well does the whole market do?

7.1. *Undifferentiated Markets—A Baseline*

Let us start with the most basic of the markets. Each of the market agents is randomly initialized with a set of goods to sell and a set of buy orders to fill.

In this run, the agents are not capable of changing their buy and sell lists, or estimating their future utility. They just blindly trade in auctions, and attempt to maximize their utility one auction at a time.

At the end of 20 auction periods, the messages were collected and collated into an adjacency Matrix and a social network. The social network diagram on figure 3 illustrates the connections that traders had to make to establish to other traders in order to complete one transaction. Even given fairly advance matchmaking tools (Sycara et al., 1999b), an agent had to communicate with as many as 12 other agents before being able to complete a transaction.

As the adjacency matrix (figure 3) shows, there is no clear pattern to which agents have to communicate in order to complete transactions, which results in higher transaction costs and necessity for more communication to achieve a result.

7.2. *Baseline Results*

The baseline results are not at all surprising, or even useful (Tables 1–3). There is a strong correlation between the average network load and market saturation—higher market saturation

Table 1. Transaction cost in baseline configuration.

	Market saturation		
Market size	0.1	0.5	0.8
25	4.7200	18.0800	23.2000
50	9.4800	36.0400	47.4000
100	18.9800	74.5600	94.9200

Table 2. Average agent utility (per transaction) in baseline configuration.

	Market saturation		
Market size	0.1	0.5	0.8
25	0.81	0.63	0.47
50	0.72	0.41	0.28
100	0.45	0.37	0.14

Table 3. Average network load (msg/sec) in baseline configuration.

	Market saturation		
Market size	0.1	0.5	0.8
25	7.2000	22.1000	28.9000
50	21.9000	92.8000	118.2000
100	92.2000	364.9000	473.4000

results in necessity for more communication to achieve a similar result. As market size and saturation increase, transaction costs go up, thus bringing down average agent utility.

The baseline experiment shows that a undifferentiated continuous double auction-based market is only viable while the number of agents involved is small and the market saturation is low. Any increase in the market saturation or market size results in rapidly increasing costs and, therefore, decreasing utility of individual agents as well as overall social welfare.

In a large and saturated market, there is a greater chance that lower-priced goods will appear, however the competition for them is very intense and many agents are forced to contend with buying higher-priced versions of the same good.

7.3. *Emergence of Market Segmentation*

As Hannan and Freeman (1977) state, "... Organizations may insure reliable performance by creating specialized units..." or-retreating into market niches that allow a highly specialized organization to thrive.

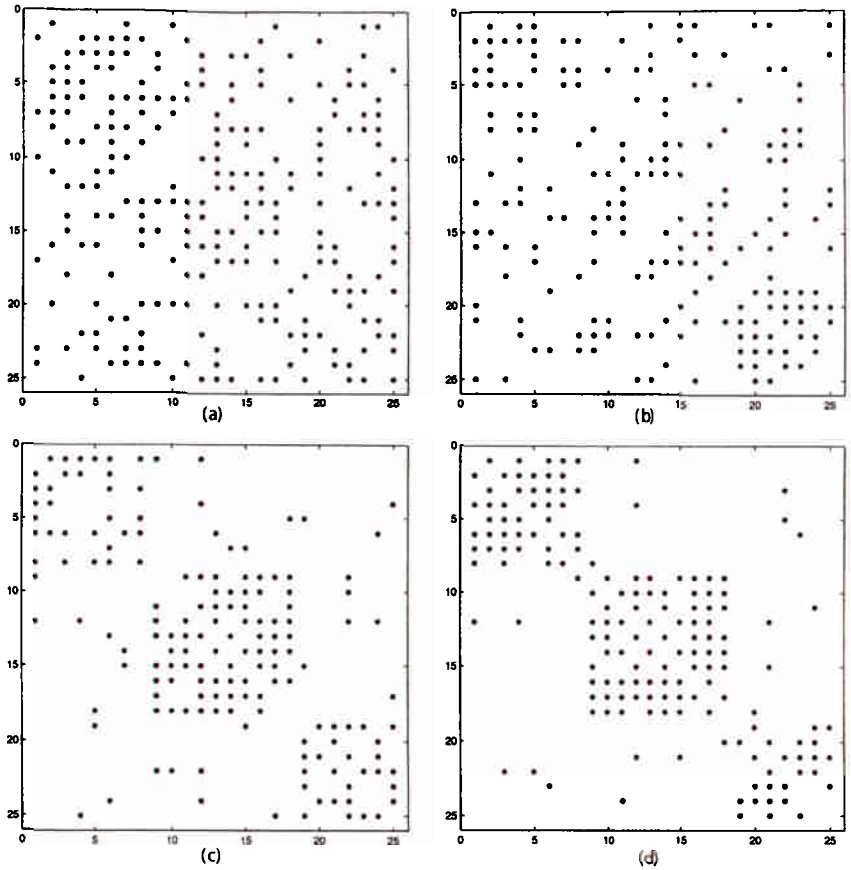

Figure 4. Emergent market segmentation, 25 agents: (a) Unsegmented market, (b) After 100 clearing periods, (c) After 250 clearing periods, (d) After 500 clearing periods.

To simulate this process, agents were allowed to add and drop goods from their lists, based on the utility they gain from the transaction—thus allowing an agent to become as much of a generalist or specialist as the market conditions allow.

As the agent does business, utility from each transaction is normalized to be in 0–1 range, and running averages for each good are kept. The probability of the agent dropping an item from its list is directly proportional to the normalized utility.

Adding items is a risky preposition. In the real world, it might be possible to make estimates of what profits other agents are making, but in this simulation this data is purely private. Thus, adding goods becomes a matter of chance. The probability of an agent adding a good to its inventory is inversely proportional to the overall normalized utility (i.e. the lower the agent's profits, the more likely it is to try a new line of business).

Table 4. Effects of market specialization.

	a	b	c	d
Network load	21.20	18.28	12.44	8.80
Transaction cost	18.2	13.41	9.3	4.2
Overall saturation	0.3280	0.3312	0.2576	0.2320
Average agent utility	0.45	0.42	0.68	0.73

The following set of adjacency matrices is the snap-shot of one of the runs of the system (figure 4). Patterns similar to these were observed in all of the 20 runs of the market simulator.

7.4. *Global Patterns from Local Behavior*

The emergent specialization has a profound effect on the market conditions (see Table 4). As agents specialize in selling one particular item, the network load decreases dramatically. As a consequence of a lower network load the transaction cost also decreases, which allows agents go get higher per-transaction utility.

The overall market saturation also decreases, thus limiting the amount of competition in each of the market sectors and virtually eliminating any cross-talk between different sectors. This does not sound like good news for the market. However it has been noted in the literature (Hannan and Freeman, 1977) that a market shakedown often occurs after initial explosion. At the end of each shakedown the number of agents in a given market sector stabilizes at the maximum number of agents that can be sustained in the sector.

8. Conclusions and Future Work

In this paper we have demonstrated a multi-agent model of a marketplace populated by self-interested adaptive agents. The model illustrates the segmentation of commodity markets by specialty—an emergent behavior borne out of local profit maximization motives. However, the local behaviors result in advancement of the global good—since the increase in segmentation of the market resulted in higher utility values, lower transaction costs and lower network loads for all agents in the market.

However, the model does not yet completely reflect patterns of interaction that occur in real markets. One of the most important aspects in terms of specialization is the advent of organizations of traders. In a real market, a trading firm employs many traders specializing in different sectors of the market. However, the utility calculation and adaptation is done at the managerial level, which is above the market segments.

It is in our immediate plans to introduce management and hiring protocols and decision-making structures into the MarketSim—which would allow agents to hire each other and create organizations. The model will then be used to study emergence of organizational structures and effects of organizational structure upon the market performance of a firm.

Due to its distributed nature, the MarketSim model has proved to be a very scalable system, accommodating large numbers of agents and dense social networks. However, there remain several bottlenecks that limit the message throughput of the agents. In particular, in larger markets the blackboard was found to have delays. In the future, these bottlenecks will be eliminated so larger studies could be done to estimate the role of scale effects in simulated markets.

Acknowledgments

This research has been supported, in part, by National Science Foundation under the IGERT program for training and research in CASOS. Additional support was also provided by the Center for Computational Analysis of Social and Organizational Systems at Carnegie Mellon. The views and conclusions contained in this document are those of the author and should not be interpreted as representing official policies, either expressed or implied of the National Science Foundation or the U.S. government.

References

Bichler, M. (1999), "Decision Analysis—A Critical Enabler for Multi-Attribute Auctions," in *Global Networked Organisations, Proceedings of the 12th International Bled Electronic Commerce Conference.*
Chemdex.com, 1999.
Decker, K., K. Sycara and M. Williamson (1996), "Matchmaking and Brokering," *Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96).*
Economou, G., M. Tsvetovat, K. Sycara and M. Paolucci (2001), "Interaction without Commitments: An Initial Approach," *Proceedings of International Conference on Autonomous Agents.*
FastParts.com, 2000.
FreeMarkets.com, 2000.
Hannan, M.T. and J. Freeman (1977), "The Population Ecology of Organizations," *American Journal of Sociology*, 82(5), 929–964.
Howley, A.H. (1950), *Human Ecology: A Theory of Community Structure*. Roland.
Jones, R.M., R. Wray, M. van Lent and J.E. Laird (1995), "Planning in the Tactical Air Domain," *Planning and Learning: On to Real Applications, Papers from the 1994 Fall Symposium (Technical Report No. FS-94-01)*, AAAI Press, Menlo Park, CA. Soar 95.27.
Smith, V. (1996), "Economics in the Laboratory," Technical Report, University of Arizona.
Sycara, K., K. Decker, A. Pannu, M. Williamson and D. Zeng (1996), "Distributed Intelligent Agents," *IEEE Expert—Intelligent Systems and Their Applications*, 11(6), 36–45.
Sycara, K., J. Lu, M. Klusch and S. Widoff (1999a), "Dynamic Service Matchmaking Among Agents in Open Information Environments," *ACM SIGMOD Record, Special Issue on Semantic Interoperability in Global Information Systems.*
Sycara, K., J. Lu, M. Klusch and S. Widoff (1999b), "Matchmaking among Heterogeneous Agents in the Internet," *Proceedings AAAI Spring Symposium on Intelligent Agents in Cyberspace.*

Maksim Tsvetovat is a Ph.D. student at Carnegie Mellon University. He has received a B.S. and M.S. degrees in Computer Science from University of Minnesota, with specialization in artificial intelligence and distributed systems. His interests include large-scale social simulation, self-organizaing systems and multi-agent systems.

Computational & Mathematical Organization Theory, 8, 235–241, 2002
© 2002 Kluwer Academic Publishers. Manufactured in The Netherlands

NETEST: Estimating a Terrorist Network's Structure—Graduate Student Best Paper Award, CASOS 2002 Conference

MATTHEW J. DOMBROSKI
Department of Engineering & Public Policy, Carnegie Mellon University, Pittsburgh, PA 15213, USA
email: mdombros@andrew.cmu.edu

KATHLEEN M. CARLEY
Institute for Software Research International, Carnegie Mellon University, Pittsburgh, PA 15213, USA
email: kathleen.carley@cmu.edu

Abstract

Since the events of September 11, 2001, the United States has found itself engaged in an unconventional and asymmetric form of warfare against elusive terrorist organizations. Defense and investigative organizations require innovative solutions that will assist them in determining the membership and structure of these organizations. Data on covert organizations are often in the form of disparate and incomplete inferences of memberships and connections between members. NETEST is a tool that combines multi-agent technology with hierarchical Bayesian inference models and biased net models to produce accurate posterior representations of a network. Bayesian inference models produce representations of a network's structure and informant accuracy by combining prior network and accuracy data with informant perceptions of a network. Biased net theory examines and captures the biases that may exist in a specific network or set of networks. Using NETEST, an investigator has the power to estimate a network's size, determine its membership and structure, determine areas of the network where data is missing, perform cost/benefit analysis of additional information, assess group level capabilities embedded in the network, and pose "what if" scenarios to destabilize a network and predict its evolution over time.

Keywords: covert networks, terrorist organizations, Bayesian inference models, biased net theory, biased networks, multi-agent systems, network estimation, social network analysis

The events of September 11, 2001 have illustrated the need for defense and investigative organizations to prepare and innovate for asymmetric and unconventional warfare in the 21st century. Our military and intelligence communities are unprepared to counter the elusive and deadly threat posed by terrorist and other covert organizations in this new century (PAM 525-5, 1994). The inter-state model of warfare has been replaced by intra-state warfare, which largely consists of guerilla and terrorist forms of warfare (Smith et al., 2001). Military and intelligence organizations in the United States are adept at utilizing training, tools, and weapons designed to counter conventional threats to national security. However, the enemies of the United States have evolved into network forms of organization that do not obey traditional borders between states and employ unconventional techniques. A networked form of organization promotes the terrorists' covert nature and decentralizes the terrorist association, allowing parts of the organization to effectively operate almost autonomously