

Congregating and Market Formation

Christopher H. Brooks and Edmund H. Durfee

Artificial Intelligence Laboratory

University of Michigan

Ann Arbor, MI 48104

{chbrooks, durfee}@umich.edu

Abstract

Agents in a multiagent system are not typically entirely self-sufficient; instead, they frequently need to enlist other agents to perform tasks for them or to exchange goods or services with them. This creates a problem: how can an agent efficiently locate other agents to work or trade with? As the number of agents grows, the cost of this computation can become prohibitively large. One solution to this is for the system to self-organize into smaller groups of agents. In this paper, we apply the idea of *congregating* to a model of an information economy. We illustrate how participants in this economy can self-organize into a set of markets such that agents are able to find suitable partners while retaining low computational costs. We show how congregating can help allocation problems scale to large populations by allowing agents to interact locally.

Introduction

In a dynamic, open multiagent system, agents have to continually make and re-evaluate their decisions as to which other agents they should interact with. These interactions might include the buying and selling of goods, negotiation over a set of tasks to perform, or cooperation to achieve a joint goal. A great deal of research in both game theory and artificial intelligence, cooperative game theory (Mas-Colell, Whinston, & Green 1995) and coalition formation (Shehory & Kraus 1998), auction theory (Wellman *et al.* 2001), and team formation (Tambe 1997) has focused on how a group of agents can make these sorts of decisions. Two problems typically arise in these sorts of approaches: scaling to domains with large numbers of agents, and dealing with the fact that the composition of the agent population is not typically known to any single agent.

Our research addresses the problem of how a large population of agents can, in a self-organizing fashion, separate itself into subgroups known as *congregations*. Ideally, each congregation should contain a collection of agents who would like to interact with each other. Once these groups exist, standard allocation algorithms such

as auctions or coalitions can be used within each congregation to decide which agents should perform particular tasks or exchange goods.

Congregations are a useful self-organizing mechanism for several reasons. First, they allow the system to scale: once an agent has joined a congregation, decisions as to whom it should interact with can (at least initially) consider only the other members of its congregation. As long as the congregation size remains relatively constant, the population can grow without impacting an agent's decision process. Second, it provides the system with a type of institutional memory. The congregation can be considered as an entity unto itself, much in the way that we can talk about an anthill or a football team both in terms of its members and as a single object. The exact population of a congregation may change as particular agents leave or arrive, and yet the group as a whole is able to maintain a long-term existence. Third, it can serve as a way of coordinating agents' decisions as to which group(s) to join. By describing a congregation in a meaningful way, a system designer can hope to attract agents of particular types or interests.

Congregating can be seen as a form of 'cooperative' learning, or group-level learning, even though the participants are all completely self-interested. The population as a whole is trying to learn how to organize itself. In most multiagent systems, the agents are playing a general-sum game, in which there are particular configurations of congregations that will lead to higher utility for a large fraction of the agent population. Therefore, discovery of these congregations can be viewed as an adaptive distributed search through the space of all possible congregations, in which each agent has a partial influence on the exploration process.

In the rest of this paper, we will discuss the role of congregations in an information economy, define more precisely what we mean by a congregation and describe it in terms of a simple model of an information economy. We will then show experimental results that suggest that congregations can improve the overall utility of the agents in a multiagent system. We will conclude with future directions for studying the formation and dynamics of congregations.

Congregating in Information Economies

Our past research (e.g. (Brooks *et al.* 1999; Brooks, Durfee, & Das 2000)) has focused on how producers in an information economy can locate consumer niches and efficiently learn the preferences of consumers in these niches. It is this first problem, that of locating a niche, that is most relevant to congregating. In this problem, two or more producers are each separately selecting a set of goods to sell and a pricing mechanism for these goods. In doing this, they separate the consumer population into one or more markets, each of which buys from one producer, plus the market consisting of consumers who choose not to buy from any producer. Each producer's problem is then to find a profitable location in price/product space, subject to the decisions of the other producers.

In this paper, we take the problem of market formation one step further, and draw directly on the congregating metaphor to examine the process of market formation in a multiagent system containing a large number of producers and consumers. Elsewhere (Brooks & Durfee 2002), we present a general model of congregations, focusing on the affinity group domain. This work showed how congregating can be difficult when the solution space is very sparse, and suggested the introduction of labelers as a signaling mechanism that served to coordinate agents' decisions. It also showed how congregating could be used to improve average payoff when agents with heterogeneous preferences played a coordination game. Our current paper applies congregating to information economies. We show how congregating can be used to stimulate producers and consumers in an information economy to self-select into separate markets so that the average profit in the system is increased. The effectiveness of this approach depends upon the number of markets available, as well as the specificity of agent preferences.

In many large-scale economies, the problem of who to buy or sell from is a significant one. Mechanisms such as auctions which pair buyers and sellers can be used to do this when the number of agents is relatively small, but if the population is large or the auction is combinatorial, both the allocation of goods and the selection of bids becomes a difficult computational problem. A solution to this is to construct several smaller auctions, each of which contains buyers and sellers with 'similar' interests.

One inspiration for this approach was the University of Michigan Digital Library (UMDL) (Durfee *et al.* 1998), which was a market in which automated agents would buy and sell information goods. The UMDL contained consumers with many different interests, and providers offering a wide set of different types of articles. Each consumer had a User Interface Agent (UIA) that located sellers of relevant information goods. The sellers were also represented by agents, known as Collection Interface Agents (CIA). These agents interacted

through auctions, each of which was described using an ontology which indicated the goods being bought and sold. These auctions were instantiated via the Auction Manager Agent (AMA), which would dynamically create new auctions when it noticed unsatisfied consumer demand and delete auctions when they experienced a lack of activity. This was a centralized process which allowed for the partitioning of the market space into separate congregations, making the location of potential customers or goods a manageable process for the agents involved, as well as the human users.

In our exploration of congregating and market formation, one goal is to reproduce this process without the intermediation of a centralized market maker (the AMA). Instead, we propose that this functionality can be handled by a distributed set of market makers, each of which has the goal of attracting buyers and sellers such that its market is successful (measured in terms of either number of trades or surplus generated). In essence, we would like for the formation of multiple markets, each of which is composed of buyers and sellers with complementary interests, to be an emergent phenomenon. The purpose of this work is to identify whether congregating can in fact be a viable strategy for a large group of self-interested agents.

Congregations

In this section, we define more precisely what we mean by a congregation, both generally and in terms of our information economy model. A more detailed presentation of this model (with minor differences) can be found in (Brooks & Durfee 2002).

We begin by considering a general model of a multiagent system. Let $A = \{a_1, a_2, \dots, a_n\}$ be a set of agents in the system at time t . (For ease of presentation, we will assume discrete time. However, nothing in our model requires it.) The essential criterion for congregating to be an interesting problem is that each agent needs to interact with some other subset of the agents in the population. In utility-theoretic terms, each agent a has a utility function $U_a : P(A) \rightarrow \mathbf{R}$. ($P(A)$ is the power set of all agents.) This function indicates the utility that agent a receives from interacting with a particular set of other agents.

It is assumed that an agent cannot interact with every other agent, due to computational and communication constraints. (Recall that n is large.) Therefore, at every time t an agent will need to choose a subset of agents to interact with.

Congregations are a way of simplifying this decision problem; rather than choosing agents from the entire population of size n , an agent can join a congregation of size $c \ll n$. This allows the agent to reduce the size of its search problem by considering the c agents in its congregation to interact with. Additionally, in domains in which agents have symmetric interests (meaning that if agent a wants to interact with agent b , then b wants to interact with a), forming congregations will improve an

agent’s chances of a ‘successful’ interaction by allowing agents to group together with other like-minded types.

Most generally, a congregation c is simply a tuple $\langle l, A_c, t \rangle$, where A_c is a set of agents $\{a_i, \dots, a_j\}$ who have all collocated at a ‘location’ l at time t . (This need not be a physical location; it could be a particular multicast address, radio frequency, or mailing list. The point is that they communicate only with other agents in the congregation.) We assume for simplicity that an agent is only a member of one congregation at a time; future work will relax this assumption. A congregation is defined *intrinsically* by its membership. It may be useful for the members of a congregation or for other agents, such as market makers, to try to describe a congregation *extrinsically* by characterizing some quality shared by its members, but this is just a label.

Let us make this all more concrete by placing it in terms of an information economy. Our set A of agents consists of two disjoint subsets: a set B of buyers and a set S of sellers.¹ Our commodity is information goods: an information good is an article of a particular category $c \in \{c_1, \dots, c_k\}$. Examples of categories include Sports, Arts, International News, and so on. At each time t , a seller is able to select a type of good to offer. Each buyer has a category c^* of good that it most prefers; articles in this category have a reservation value r . A good from another category c' is valued at:

$$V(c') = r \left(1 - \frac{|c^* - c'|}{k} \right) \quad (1)$$

Note that this assumes that categories are arranged according to a ‘similarity’ metric on the k axis. While this is a simplification, it is sufficient for our needs. It also provides our model with some similarities to the Hotelling model, which is a common economic framework for product differentiation. Hotelling models typically contain one or more product dimensions along which producers can differentiate themselves, and consumers with heterogeneous preferences along these dimensions. Each producer must determine where to locate, given the expected locations of the other producers. Anderson, et al. (Anderson, de Palma, & Thisse 1992) provide a thorough overview of Hotelling-style models.

In our model, each agent a has a threshold τ_a which indicates the fraction of r it wants to receive from a transaction. This threshold can be interpreted either as a production or consumption cost, or as a form of satisficing, whereby all transactions yielding utility greater than $\tau_a \times r$ are considered satisfactory by agent a .

In a simple world, this would be all that is needed for a model; buyers and sellers could be paired up using some sort of predetermined mechanism, either a decentralized mechanism such as Contract Net (Davis & Smith 1983), in which agents broadcast offers to each

other, or a more centralized mechanism such as an auction, where bids are submitted to a central authority which computes allocations. In either case, there is some computation that must be performed in order to determine an allocation. In the case of Contract Net, this computation is the exchange of messages, and in the case of an auction, it is the winner determination algorithm. In either case, this computation does not come for free; we assume that the agents in the system must pay for it. By separating into separate congregations, they can ease their computational burden, hopefully without a large decrease in the efficiency of the allocation of goods. In fact, the goal of the experiments in this paper is to determine how much gross efficiency is lost, how much computation is saved, and the resulting net costs as the number of congregations is varied.

In order to capture this notion of agents separating into subgroups, our model also consists of a set $M = \{m_1, \dots, m_q\}$ of markets. These are the locations described above; buyers and sellers will congregate here. Each buyers and seller will choose exactly one market from M to join. At a time t , the market will close and an allocation will be computed for all agents in the market. The cost of this allocation will be shared equally amongst all agents in the market. The mechanism we use is a standard N th-price Vickrey auction. This was chosen for simplicity; since all goods are privately valued and there is no resale, the dominant strategy is for agents to bid their actual valuations for articles (MacKie-Mason & Varian 1994). However, any allocation mechanism can be used in this model. If costs are superlinear in the number of agents, the same qualitative effects will be seen. Each buyer bids its reservation value and the prices and allocation of goods is computed. Each agent can then decide whether to remain in this congregation or leave for another; this decision will be based on whether the value received (less computation cost) exceeds the agent’s threshold.

Clearly, in a single market, the allocation which maximizes efficiency can be computed. This will be used as a benchmark in our experiments. However, computing the optimal allocation in a combinatorial auction is NP-complete (Sandholm 2001), which means that the computational costs will be prohibitively large as the number of agents in a market becomes large. Even in simpler allocation problems, such as coalition formation, the number of messages exchanged is polynomial in the number of agents in the system, which leads to scalability problems as the number of agents becomes large.

Elsewhere (Brooks & Durfee 2002), we have analyzed the convergence properties for a simple congregating model, focusing on the length of time needed for the optimal set of congregations to form. That work showed that congregating can lead to higher net utility in domains where agents try to locate other agents of the same ‘type’. One observation from that work is that congregating is useful because it helps to build a ‘crit-

¹We assume that there is no resale; buyers consume their goods immediately. Economies containing supply chains are also a promising area of application for congregations, but are not treated here.

ical mass;” a group of agents will form a congregation that is mutually beneficial, and so the congregation will persist over time. This provides a focal point for the attraction of other agents that may want to be a part of this congregation. Essentially, the congregation allows a subset of agents to “hold still” and allow others to find them.

In this paper, we are explicitly interested in the learning dynamics of market formation: what do the transitional congregations look like, how are agents’ aggregate profits affected by the congregating process, and whether congregating can improve overall profit by providing this sort of critical mass. Due to the complexity of the problem and our desire to examine the transitional behavior of the system, we focus on experimental methods to determine when congregating is a viable strategy, and how overall profits (as a measure of system performance) change as the number of markets is altered.

Using Congregations to Improve Net Profit

In this section, we present experiments demonstrating the usefulness of congregating when agents must incur a computation cost for computing an allocation of goods. We show how congregating is a useful strategy when agents must pay for computation, and study the efficiency of the allocation (both gross, pre-computation cost and net, post-computation) as the number of markets is varied.

In the first experiment, we consider an information economy composed of 50 producers and 50 consumers. There are 10 categories of information goods; at the beginning of the experiment, each producer randomly chooses a category of good to sell. Each consumer also has a favorite category (drawn from a uniform distribution) and preferences over non-favorite goods indicated by equation 1. Reservation values for a consumer’s most-preferred good are drawn from $U[5, 10]$. The number of markets is fixed at the beginning of the experiment; it is assumed that all agents know m , the number of markets, and all agents have agreed to share equally in the costs of computation in determining an allocation. We assume that computation has a cost of 0.1 for each message sent or comparison made. Since we are using a generalized Vickrey auction to compute allocations, it is straightforward for the market to hold an auction, determine the cost of this computation, and charge each agent accordingly.

After each market closing, every agent is able to change markets. If the producer receives less than $\tau \times r$ for its good, or if a consumer receives a good that it values at less than $\tau \times r$, it will move to a new congregation chosen at random. In our initial experiments, τ was set to 0.5 for all agents. This is varied in section . (Note that we are not advocating random movement as an optimal agent strategy. Instead, we are interested in characterizing the behavior of the system as the number

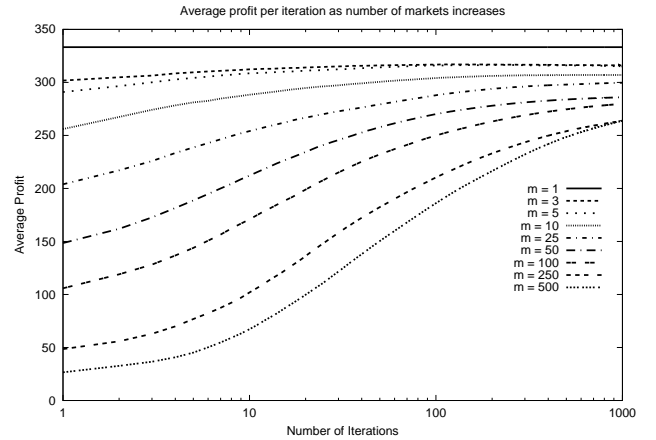


Figure 1: Cumulative profit per iteration over all sellers, averaged across 10 runs of 1000 iterations each (x axis is log scale)

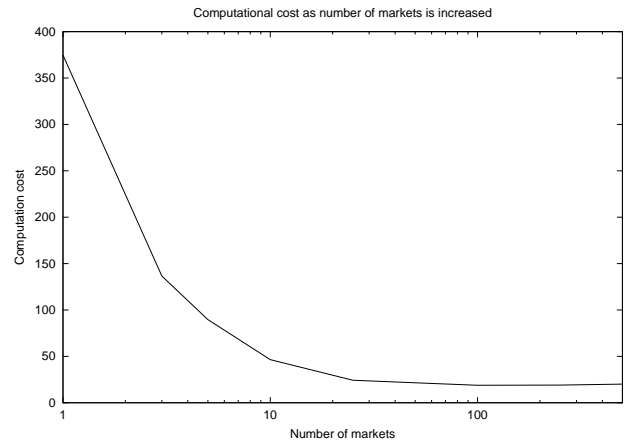


Figure 2: Cost of determining the optimal allocation of goods as the number of markets increases (averaged over 50 instantiations with 50 producers and consumers).

of congregations is varied when agents exhibit some sort of adaptive behavior. Random movement is meant to provide an approximation of “typical” agent behavior when the population is large.)

Figure 1 shows the average cumulative profit achieved by all sellers (averaged across 1000 iterations) as the number of markets is increased. As the number of markets is increased, initial profits are much lower, as it becomes more difficult for each agent to locate a congregation in which it can find a suitable match. Also, the system as a whole reaches a suboptimal configuration. However, this figure does not take into account the costs of computation. Figure 2 shows how computational costs fall off quickly when the economy contains more than one market. These two graphs are combined in figure 3. In this figure, we can see that, even though a

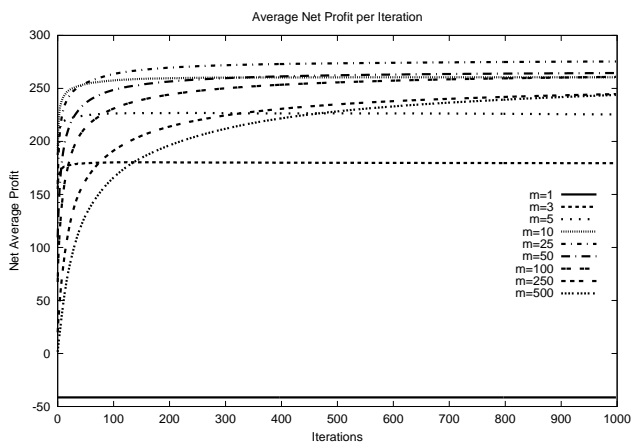


Figure 3: Net cumulative profit per iteration over all sellers, averaged across 10 runs of 1000 iterations each.

single market is able to compute the optimal allocation, the computational costs are high enough that producers end up with a negative net profit. As the number of markets is increased, the computational cost falls off much more quickly than profit, until at 25 markets the average cumulative net profit is maximized. As more markets are added, net profits begin to decline, as the computational savings is outweighed by the difficulty of finding a market containing a suitable matchup.

In figure 3, we can see the dynamics of market formation as the number of markets is increased. With a very small number of markets (3 or 5) a stable configuration is reached almost immediately. As we increase the number of markets to 10 or 25, convergence is slightly slower, but the configuration of congregations that is found yields higher average profit. As the number of markets is further increased, the quality of the configuration falls off slightly, and convergence time becomes much longer. It is difficult to quantify this comparison directly from figure 3, due to the number of lines and the fact that total performance is actually the area under each curve. Figure 4 makes this tradeoff more explicit; it shows the average net profit (over 1000 iterations) as the number of markets is increased. Here we can clearly see that net profit increases steeply as the number of markets is increased, and then falls off after $m = 25$.

We can also see that introducing too many markets, as in the case of $m = 250$ or $m = 500$, introduces too much inefficiency without a significant extra reduction in computational cost. Buyers and sellers spend a great deal of the initial iterations in a market with no suitable agents.

In our next experiment, we examine whether the advantages of congregating scale as the number of agents in the system is increased. We vary the number of producers and consumers from 10 to 500 of each, in each case setting the number of markets to half the number

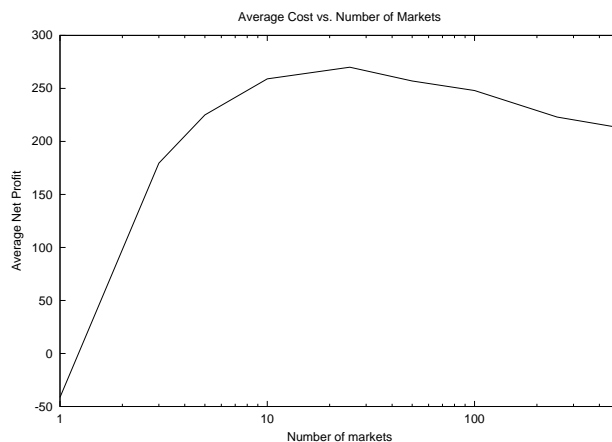


Figure 4: Average net profit (over 10 runs of 1000 iterations each) with 50 producers and 50 consumers as the number of markets is varied. (x axis is log scale)

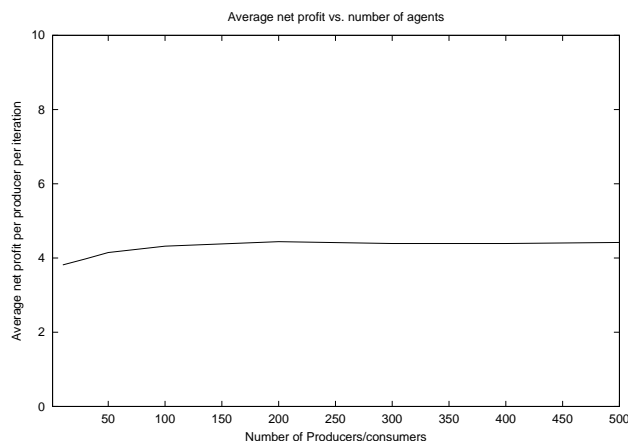


Figure 5: Net profit per iteration per producer as the number of producers (and consumers) is varied between 10 and 500.

of producers, and compare the net (post-computation costs) per producer. The average profit per producer over the course of the run is presented in figure 5.

As we can see, congregating is in fact a method that is able to scale to large numbers of agents. Profits are invariant as the number of agents in the system is increased. (In fact, net profits for small numbers of agents are slightly lower; this is an artifact of the way in which agents are randomly generated. For low numbers of agents, it is less likely, for any given producer, that there is a consumer whose most desired good corresponds to that produced.) As long as the system designer has a rough idea of how many other agents are in the economy and is able to construct the appropriate number of markets (half the number of producers in this case), market size remains constant, the economy can scale,




Figure 6: An illustration of the changes being performed on a simplified landscape. The top figure shows a simple two-peaked profit landscape. As agents’ τ is increased, the “sea level” increases, leading to the second figure. The original landscape is indicated by a dotted line. When consumers reduce the number of categories they value, the result is the attenuated landscape shown at the bottom. The heights of the peaks remain constant, but the area of the plateaus increase.

and producers and consumers are still able to find each other and make successful transactions.

Of course, these results are dependent on the particular costs and reservation values chosen. The higher the ratio of reservation values to computation costs, the more utility will be placed on an optimal solution. On the other hand, as computation costs become a significant portion of net profits, agents will prefer cheaper solutions and congregating will gain appeal.

Varying Consumer Preferences

A large part of the effectiveness of congregating comes from its ability to serve as a coordination mechanism for agents’ decisions. A subset of agents will move into a particular market and, if they are happy, stay there. This reduces the complexity of the problem for agents who have yet to find a suitable market, since fewer agents are moving. If we visualize the congregating process as a search over a landscape, congregating allows the search process to move along a gradient more effectively. Of course, this assumes that the landscape has a gradient. In this section, we alter the gradient of the landscape and examine the effectiveness of congregating.

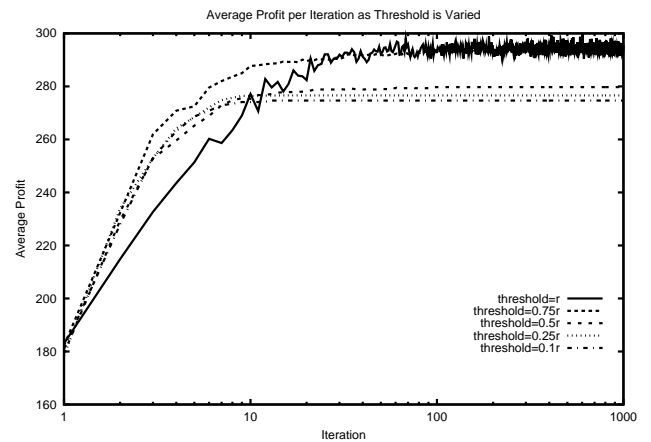


Figure 7: Net profit per iteration as the threshold is varied between 0.1 and 1, averaged over 10 experiments. r is the agent’s reservation value. X-axis is log scale.

In our first experiment, we alter agents’ threshold τ . Recall that τ is the fraction of reservation value an agent must receive to be satisfied with a transaction. As τ increases, an agent is happier in fewer congregations. In terms of the search landscape, this lowers all the peaks equally, as if there was a flood. Fewer points on the landscape have positive profit, but the topology is not deformed. Figure 6 illustrates this.

Experiments were again conducted with 50 consumers and 50 producers. All experiments were performed with 25 markets, seen to be the optimal number in our previous experiments. τ was varied between 0.1 and 1.0. Results are shown in figure 7

As we see from figure 7, when $\tau = 1$, the process takes longer to converge to a solution, but the quality of the final solution is improved. When τ is small, the process converges more quickly, but to a local optimum. The encouraging result here is that even when the threshold is high, meaning that agents have a high cost of producing or consuming articles, congregating still quickly finds a desirable solution.

In our second experiment regarding consumer preferences, we alter the *number* of categories each consumer values positively. We do this by giving each consumer a value n which indicates the number of categories they value positively. If an article offered by a producer is of a category less than $\frac{n}{2}$ away from the consumer’s preferred category, it is valued according to 1. Otherwise, it has value 0. This has the effect of attenuating the congregating search space. The gradient around an optimum is increased, producing steeper peaks and larger flat areas. While the profits attained by the system at the optima do not change, the profits for non-optimal configurations of congregations decrease, since consumers value fewer categories. In essence, the search problem becomes more difficult. This is also illustrated in figure 6.

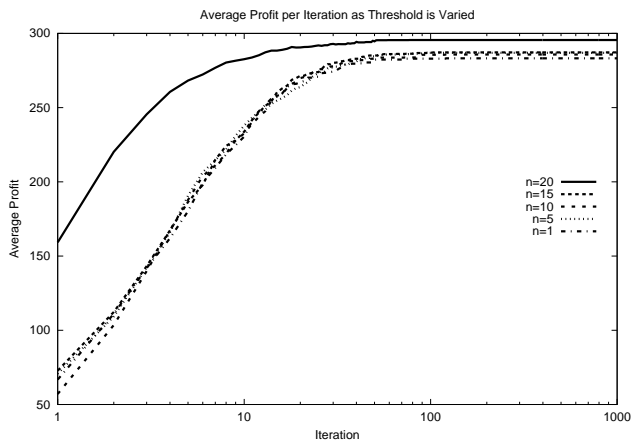


Figure 8: Net profit per iteration per producer as the number of categories (out of a possible 20) valued by each consumer is varied. Results are averaged over 10 experiments. X-axis is log scale.

Once again, experiments were performed with 50 consumers, 50 producers, 25 markets and 20 potential categories. τ was fixed at 0.5. Results for this experiment are shown in figure 8.

As we can see from figure 8, as consumer preferences become more specific, the search process becomes more difficult; many configurations of congregations do not contain a critical mass of agents receiving a positive profit. In terms of the search landscape, the size of the plateaus increases. The more interesting result is that, even though the search process becomes more difficult, through the use of congregations, the agents in the system are still able to self-organize into congregations which yield a large fraction of the optimal profit. This tells us that congregating is a method which is robust to the particular shape of agent preferences; even when consumers have very specific preferences as to the goods they want, congregating allows producers and consumers to locate each other.

Discussion

In this paper, we have provided an example of how congregating can be applied to information economies and used to model the process of market formation. As we have seen, if the participants in a market must pay for their computation, they are often happier to realize an allocation which yields slightly less utility, but can be found at a much lower computational cost. The mechanism is essentially a self-organizing one; initially, agents which are happy in a market will stay there, providing a fixed location that other ‘like-minded’ agents can find. The only common knowledge that is required is that the agents all agree on the existence of a set of markets at the beginning of the economy’s lifetime. Once this initial commitment is made, an agent need not consider the global state or worry about the identity of all the

other agents in the system; it can simply concern itself with the other agents in its congregation.

We have also shown that congregating is a technique which allows multiagent systems to scale up to large numbers of agents. When agents must consider the costs of computation in determining who to purchase from or how to allocate goods, congregating allows the system to scale so that net profit remains relatively constant as the system grows.

Additionally, we have shown that congregating is relatively robust to the shape of agent preferences. Even when consumers have a very high processing cost, or very narrow preferences, congregating still allows producers and consumers to make satisfactory exchanges, although the time needed to discover the optimal configuration of congregations increases.

One question that must be asked is the applicability of these results to other domains. Our experiments rely on the assumptions that agents have utility functions which they are able to evaluate exactly, that agents know their preferences, and that agents have some preference over who they choose to associate with. Many economic domains meet these assumptions, but there are other domains in which agents cannot easily express their utility functions. In addition, we have not touched upon the problem of market formation and advertising, apart from declaring that markets exist and that all agents know about them. A crucial problem in the UMDL was the description and dissemination of information about these markets, so as to encourage the congregating of the ‘right’ agents. This remains a topic for future research.

There are many other potential directions for this research. As noted previously, goods in this economy are immediately consumed. If goods are instead transformed and resold, then the need for well-constructed markets becomes even more important. Additionally, this work does not explicitly treat nonstationarity in the consumer population. An important question to ask would be whether the existence of congregations makes consumer and producer entry more difficult.

Finally, we earlier stressed the notion of a congregation as an entity unto itself. This leads one to consider notions of group selection, whereby it becomes rational for agents within a congregation to behave altruistically, so as to increase the fitness of the rest of the congregation and improve their chances when competing with agents outside of the congregation. In environments in which agents typically interact within a congregation, but occasionally encounter agents from the population at large, this can be a very effective strategy (see (Wilson 1980) for ecological examples). This would also lead us to examine the relationship between individual and group-level learning: as individual agents develop more sophisticated strategies, how does the composition of the congregation change?

Acknowledgments

The authors would like to thank Jeff MacKie-Mason and Bob Gazzale for their contributions to this research. This work was supported in part by NSF grants IIS-9872057 and IIS-0112669.

References

- Anderson, S. P.; de Palma, A.; and Thisse, J.-F. 1992. *Discrete Choice Theory of Product Differentiation*. Cambridge, Massachusetts: MIT Press.
- Brooks, C. H., and Durfee, E. H. 2002. Congregation formation in multiagent systems. *Journal of Autonomous Agents and Multiagent Systems*. To appear.
- Brooks, C. H.; Fay, S.; Das, R.; MacKie-Mason, J. K.; Kephart, J. O.; and Durfee, E. H. 1999. Automated strategy searches in an electronic goods market: Learning and complex price schedules. In *Proceedings of ACM EC-99*, 31–40.
- Brooks, C. H.; Durfee, E. H.; and Das, R. 2000. Price wars and niche discovery in an information economy. In *Proceedings of ACM Conference on Electronic Commerce (EC-00)*, 95–106.
- Davis, R., and Smith, R. G. 1983. Negotiation as a metaphor for distributed problem solving. *Artificial Intelligence* 20:63–109.
- Durfee, E. H.; Mullen, T.; Park, S.; Weinstein, P.; and Vidal, J. 1998. The dynamics of the UMDL service market society. In Klusch, M., and Weiß, G., eds., *Cooperative Agents II*, Lecture Notes in Artificial Intelligence. Springer. 55–78.
- MacKie-Mason, J. K., and Varian, H. 1994. Generalized Vickrey auctions. Technical report, University of Michigan. available at: <http://www-personal.umich.edu/jmm/papers/gva3.pdf>.
- Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*. New York: Oxford University Press.
- Sandholm, T. 2001. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*. to appear.
- Shehory, O., and Kraus, S. 1998. Methods for task allocation via agent coalition formation. *Artificial Intelligence* 101:165–200.
- Tambe, M. 1997. Toward flexible teamwork. *Journal of Artificial Intelligence Research* 7:83–124.
- Wellman, M. P.; Walsh, W. E.; Wurman, P. R.; and MacKie-Mason, J. K. 2001. Auction protocols for decentralized scheduling. *Games and Economic Behavior* 35(1/2):271–303.
- Wilson, D. S. 1980. *Natural Selection of Populations and Communities*. Menlo Park, California: Benjamin/Cummings.