An Introduction to Congregating in Multiagent Systems

Christopher H. Brooks Edmund H. Durfee Aaron Armstrong*
Artificial Intelligence Laboratory
University of Michigan
Ann Arbor, MI 48109
{chbrooks,durfee,armst}@umich.edu

Abstract

We present congregating *both as a metaphor for describing and modeling multiagent systems (MAS) and as a means for reducing coordination costs. We show how congregations can be used to explain and predict the behavior of self-interested agents that are searching for other agents to interact with. This framework is integrated with Vidal and Durfee's CLRI framework [11] for evaluating learning within MAS. We provide experimental and analytical results which describe how the difficulty of the congregating problem increases exponentially with the number of agents, and present a solution to this in the form of* labelers*, which are agents that assign a description to a congregation, thereby reducing agents' search problem.*

1 Introduction

In a multiagent system, self-interested agents must often decide which other agents they want to interact with. The nature of these interactions may vary; perhaps they wish to buy and sell goods, exchange information about their environment, group together in order to exploit scaling effects, or simply benefit from the presence of other agents. These interactions are what makes a society more than just a collection of agents which happen to be in the same location; each agent's reward is dependent upon the agents that it interacts with.

However, as the number of agents in a system increases, the number of potential interactions that a particular agent must consider grows exponentially, since agents must potentially consider all groups of agents which they could be a part of. Something is needed to both simplify an agent's decision regarding which other agents to deal with and allow it to devote some initial energy to making this decision so as to yield more efficient interactions in subsequent iterations.

*This work was supported in part by an IBM University Partnership Grant and by the National Science Foundation under grant IIS-9872057

One way in which human societies have dealt with this problem is through the establishment of *congregations*. Human congregations include clubs, churches, marketplaces, university departments, and USENET newsgroups. Members of these congregations have devoted some upfront cost to organizing and describing themselves so that they can both reap the long-term benefits of interacting without coordinating further and also attract new agents whom they would be likely to want to interact with.

Congregations were previously described in [1]. In this paper, we continue to explore the notion of groups of self-interested agents trying to organize themselves. Much of the previous work on multiagent coordination, such as RETSINA [9], along with work on coalition formation [8] and team formation [10], as well as work on game-theoretic negotiation [6], has assumed that agents have well defined roles within an organization or specific tasks which they perform. In contrast, a congregation is a more general structure. The distinction is that congregating agents expect to have a long lifetime, during which they take on different roles, perform different tasks, and interact with different agents. By devoting some initial effort to constructing an congregation in which an agent can easily locate other agents with desirable characteristics, a community of agents can then take advantage of this structure to avoid devoting future resources to coordination.

We present congregating as a multiagent learning problem; multiple agents are learning at the same time, and the global state of the system is an important factor in determining the utility of each individual agent. The notion of coordination as multiagent learning has also been discussed in [7], where cooperative agents learned to coordinate actions to achieve a joint goal, and in [4], where self-interested agents are able to each learn a Nash equilibrium strategy. A primary difference between those works and this paper is that the congregating problem focuses on a decision as to *who* to interact with, rather than what action to choose.

We begin with a formal description of congregations and the associated learning problem of finding the correct congregation to join. We relate our model of congregation formation to Vidal and Durfee's CLRI [11] model for analyzing multiagent learning behavior. We then present experiments which analyze the difficulty of the congregating problem as the number of agents grows, and show how introducing a set of agents which label congregations can help to ameliorate this problem.

2 What is a Congregration?

In this section we define a congregation in more detail. We present some characteristics that are essential for a multiagent problem to be considered as a congregating problem, and then introduce a formal model of congregating.

2.1 Characteristics of congregations

We begin by presenting components which are essential for a multiagent problem to be considered within the congregating framework.

- Individual rationality. Each agent is assumed to have its own utility function. Agents will act solely to maximize their long-term utility, where "long-term" indicates that an agent will take a discounted estimate of future rewards into consideration when deciding with whom to congregate.

  Note that we do not require (nor do we expect to use) any notion of "group rationality." Groups of agents which receive a single lump payment as a result of the group's performance do not fit into our definition of congregations; these are more accurately described by existing work on coalition [8] or team [10] formation.

- Agents may voluntarily join or leave congregations. An essential facet of the congregating problem is that agents are free to join (or leave, or refuse to join or leave) a congregation at any time they wish. It is this decision problem (which congregation, set of congregations, or series of congregations should an agent join?) that is at the heart of this work.

- Another principal idea behind congregations is that an agent's satisfaction with a congregation is dependent upon the other members of the congregation. Since agents congregate in order to satisfy needs which they cannot satisfy alone, it seems reasonable to assume that their satisfaction will depend upon how well these needs are met. If agents are heterogeneous in their ability to satisfy these needs then agents will prefer to congregate with those agents which better satisfy their needs over those who do not.

- Agents will have repeated interaction and long-term existence. As noted above, the whole point of developing a congregation is to allow an agent to devote fewer resources making future decisions regarding who it should interact with. If an agent is making a one-shot decision, there is no value to exploring and learning information to use in future encounters.

- We also assume that agents must expend energy or pay a cost to search for suitable partners to interact with and to advertise their presence to other agents. This cost can vary depending upon (for example) the distance between agents, the number of agents reached, or the complexity of the message. Advertising messages may also take time to propagate through the system. If an agent is able to costlessly and instantaneously search through the space of all other agents to find suitable partners, then there is no need for it to spend any effort on forming congregations, since the primary function of a congregation is to reduce the long-term search and advertising cost.

2.2 Examples and Motivation

The congregating problem occurs frequently in multiagent interactions, particularly in open systems such as the Internet, with large numbers of agents that are continually arriving and leaving. An agent will find that it needs to repeatedly interact with other agents in order to satisfy its goals. It must continually make a choice as to whether to interact with previously known agents or explore further and try to find more suitable partners. One difficulty is that all of the other agents in the system are faced with the same problem; if an agent moves to a new location in search of a compatible agent at the same time that the other agent moves, they may wind up missing each other.

A congregation is a way of avoiding this search. By devoting initial resources to collocating with a desirable set of agents, an agent can then avoid having to search in future iterations. As always, there are tradeoffs to be considered. If the agent population is changing rapidly, then it may not be useful to expend a great deal of effort in developing a congregation.

Congregations bear some similarity to coalitions. In [8], a coalition is defined as a group of agents which have gathered together to either achieve a group task or allocate globally-assigned tasks to individual agents. A congregation is a more general grouping; there need not be a global task or motivation connecting each of the agents in a congregation. Agent A may be in a congregation because of the presence of B, who is in the congregation because of the presence of C, and so on. The key is that it is easier for these agents to locate a suitable 'partner' to interact with within the congregation than without.

The difficulty with congregating is that, as the number of agents and congregating places (known as loci) increase, it becomes harder for agents which should be together to find each other. Human congregations solve this problem by attaching labels with semantic content, such as 'Elks Club' or 'Methodist Church' or 'Farmer's Market' to these congregations. This allows agents with a shared vocabulary to make reasonable assumptions as to the types of agents that will join this congregation. However, a new problem is introduced: that of selecting an appropriate label. A label must be specific enough to distinguish a congregation from other congregations, yet general enough to attract the 'right' group of agents.

There are numerous examples of problems that can be described within the congregating framework. Our current research involves the structure and formation of congregations within an information economy. In this case, the agents that are congregating are consumers which have different preferences over different types of articles. A consumer would like to find a market in which articles that it prefers are sold. Producers of information goods decide what to offer, thereby acting as labelers. By offering different articles, a producer will attract different groups of consumers. Likewise, different compositions of consumers will produce different sets of aggregated preferences, thereby changing the profitability of different labels for a producer. Producers must make decisions as to how to describe and price their products; by doing so, they will attract some consumers to their congregation and discourage others. The selection of articles offered by other producers will also influence the choice of articles a producer offers; if two producers attempt to attract the same congregation, they may be worse off than if they attract separate congregations. In addition, a producer is constantly weighing the decision to incur a cost by changing its offerings or advertising more (in an attempt to find a better congregation) against the potential long-term benefits of this decision.

2.3 A Formal Model of Congregations

In this section we provide a formal model of congregating and the congregating process. This will later be connected to the CLRI model to analyze the difficulty of the congregating problem as characteristics of the problem change.

We begin with a set $Ag = \{ag_1, ag_2, ..., ag_n\}$ of agents. This set is composed of two subsets; a set $CA = \{ca_1, ca_2, ..., ca_m\}$ of *congregators* and a set $LA = \{la_1, la_2, ..., la_k\}$ of *labelers*. Each agent must be a labeler or a congregator, but not both; that is, $CA \cup LA = Ag$ and $CA \cap LA = \varnothing$.

Also, consider a set $\alpha = \{\alpha_0, \alpha_1, ..., \alpha_n\}$ of *loci*. A locus is any place in which a congregation can form. Formally, a locus has a label (defined below) and a unique name. Labelers may assign labels to a locus in order to attract a particular set of congregators.

Next, let $\Lambda = \{\lambda_0, \lambda_1, \lambda_2, ..., \lambda_i\}$ be a set of labels. These labels are generated by labelers and placed on loci. They are then used by congregators to identify some distinguishing features of a congregation. The key is that agents place a semantic meaning on each label. This meaning may be as simple as a summary of the current members of a congregation, or as complex as a description of the features that would be common to the ideal congregation that a labeler wishes to attract. We will use λ_0 to denote the "null label", which is a label that carries no information.

Finally, we define a congregation c as a locus, a set of zero or more congregators, and a set of zero or more labelers. This is, $\forall c \in C, c = \{\alpha_i, ca, la\}$, $\alpha_i \in \alpha, ca' \in P(CA), la \in P(LA)$, where C is the set of congregations that can exist in a system.

2.3.1 Congregators

A congregator wishes to maximize its long-term utility. The underlying assumption here is that a congregator will have many chances to join congregations and therefore wishes to maximize its total satisfaction over all these encounters. Its predictions of future reward may be discounted so as to reflect either uncertainty of belief or the importance of immediate rewards. Our conjecture is that a congregator can benefit in the long term by devoting some initial resources to finding a good congregation. Thereafter, its decision as to who to interact with is simplified, since suitable partners are occupying the same congregation. For this work, we assume that a congregator can only join one congregation at a time.

Each congregator has a payoff function P which maps from a congregation c to a real number: $P : C \rightarrow \mathbb{R}$. Since a congregation's payoff to an agent is dependent upon the other agents that are a part of the congregation, a congregator will not be able to fully evaluate a congregation's utility before deciding whether to join it. Only after it (and all other agents) make this decision will it be able to evaluate P. While making their decision, congregators will rely on an estimate of their payoff (denoted $\hat{P}$) which maps from loci (and the labels placed on them) to the reals: $\hat{P} : \alpha \rightarrow \mathbb{R}$.

For this work, we will assume that $\hat{P}$ is reasonably accurate, so that we do not need to worry about congregators having to learn their preferences over labels. That is, when offered a choice between two labels, congregators already know which one they would prefer. This reduces the congregators' problem to that of selecting the best congregation to join from those which exist.

Let Λ_i be the set of labels which are offered to congregators (by labelers) in iteration i. A congregator c_j's decision function δ_{c_j} is then to choose a locus which maximizes $\hat{P_{c_j}}$.

2.3.2 Labelers

A labeler's problem is potentially more complicated. As with the congregators, a labeler wishes to maximize its long-term utility. It has one decision to make on every iteration: which label to offer. There are two confounding factors which make this problem difficult: a labeler does not necessarily know the congregators' preferences over congregations (or labels), and so must learn them; and the labeler's payoff for selecting a particular label may also be contingent on the labels offered by other labelers.

Formally, a labeler la has a payoff function P which is a function of all labelers' decision functions. (Recall that labelers each decide on a label simultaneously.) $P : \lambda_{la} \times \lambda_2 \times ... \times \lambda_n \rightarrow \mathbb{R}$. As with the congregators, labelers most likely do not have access to this function, and may find it easier to instead work with a function Con which predicts the congregation that will result from a given labeling decision. $Con_{\lambda_i} : \lambda_1 \times \lambda_2 \times ... \times \lambda_n \rightarrow P(CA)$ This function tells, for a labeler's label choice, along with all other choices, which agents will join its congregation.

This function may also need to be learned over time, meaning that a labeler would have functions $\hat{P}$ or $\hat{Con}$ or both. In order to do so, it may wish to model either the strategies and knowledge of other labelers or the preferences of congregators, or both. In this paper, we shall assume that labelers are 0-level learners in the RMM sense [3], meaning that they do not explicitly model other agents and their learning processes. Future work will examine the gains that labelers can make by attempting to infer something about the learning processes of both congregators and other labelers.

3 Learning in Affinity Groups

In [11], Vidal and Durfee present a framework known as CLRI for modeling multiagent learning and predicting the behavior of multiagent learning systems. We briefly summarize this and then show how it can be applied to a restricted version of the model described above in a particular setting, namely an affinity group domain. An affinity group is a set of agents which all share some trait or preference, such as hair color or a desire to discuss LISP, and want to congregate with other agents which also have this trait and avoid agents which do not have this trait or preference. Variants of this problem occur in real life in places where agents (human or artificial) are attempting to find other agents with similar interests, such as in matchmaker systems or newsgroups.

Note that in our use of the term *learning*, we are not referring to single-agent learning (the agents portrayed within are not terribly clever), but to multi-agent learning, in the sense that we are interested in the performance of the system as a whole improving over time.

3.1 The CLRI Framework

The CLRI framework is used to model the problem of agents learning a moving target function; namely, a function that is a best response to other agents' actions, where their actions change over time. In every iteration, agents choose an action a from a set A of possible actions. Agents have a decision function $\delta^t(w)$ which returns an action for a given world state w at time t. They would like for this function to match $\Delta^t(w)$, which maps to the true optimal action at time t. In the above framework, the set of actions are congregations to join, and so this δ is identical to the δ described in the Congregators section above.

One point to note here is that the CLRI framework is based upon traditional PAC-learning assumptions [5]. In particular, it assumes that, for a given world state, there is one correct action and all the others are incorrect. There are also other assumptions which will be discussed as they become evident. We will maintain these assumptions throughout this paper. In the end, we will discuss the limitations that CLRI's assumptions place on our model of congregations and present some ideas for extending the CLRI model.

CLRI begins with two simple parameters: change rate (c) and learning rate (l). Change rate is the probability that an agent will adjust an incorrect mapping of $\delta(w)$ from time t to time $t + 1$, and learning rate is the probability that the agent will change the mapping of $\delta^{t+1}(w)$ to be equal to $\Delta^t(w)$. (That is, the agent will know what it should have done at time t.) In the congregation framework, change rate is the probability that an agent will choose a different congregation than it had previously when faced with the same choices, and learning rate is the probability that an agent will be able to say after the fact where it should have gone (and thus make the right choice the next time it is faced with this set of choices).

CLRI next considers the retention rate r, which is the probability that an agent retains a correct mapping ($\delta^t(w) = \Delta^t(w)$). In terms of congregating, this is the probability that an agent which chose the correct congregation will make this same choice again when faced with the same set of alternatives, even if it has changed its mapping of δ for other alternatives.

The fourth parameter used in CLRI is volatility (v), which is the probability that the function Δ which the agent is trying to learn changes. In congregating terms, this is the probability that a locus α which was chosen at time t from a given set of loci is no longer the correct choice at time $t + 1$ when given the same set of loci to choose from.

Since volatility can be difficult to determine through observation, CLRI provides a way to derive it using a fifth parameter: impact (I_{ij}). Impact is the effect that agent i's learning has on agent j's target function. In particular, it is the probability that agent j's Δ changes between t and $t + 1$ as a result of agent i's δ changing. In terms of congregating, this is the probability that when congregator c_i decides to choose locus α' rather than α, congregator c_j's best choice of congregations changes.

In the following sections, we analyze the difficulty of the congregating problem when no labelers are present and then show how introducing labelers can help to reduce the complexity.

4 Congregating without Information

In this first scenario, we examine the case in which there are no labelers. This means that, for all loci, the associated label is λ_0, the null label. Congregators therefore can only guess randomly as to which congregation is best for them.

We begin with the simplest case, in which there is only one affinity group and a number of loci. Once we have established some analytic results for this problem, we extend this to a model with multiple affinity groups and examine how this affects the problem complexity.

4.1 One Affinity Group

If there is only one affinity group, then all congregators simply want to congregate in the same location. The problem, of course, is that they don't know ahead of time which locus to choose. This is a very simple problem, but it lets us establish some notation and lead into scenarios with multiple affinity groups.

Assume that there are ϕ agents and m loci. Each agent a receives the following payoff:

$$Payoff(a) = \begin{cases} 1 & \text{if } |C(a)| = \phi \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where $|C(a)|$ indicates the size of the congregation agent a is in. Congregators will search randomly until they find themselves in a congregation with every other agent.

The first question we can ask is how long it will take for a set of ϕ congregators to find each other. The probability that all agents will choose the same locus on any given time step is[1]:

$\epsilon = \left(\frac{1}{m}\right)^{\phi-1}$.

Therefore, the agents will succeed in finding each other with probability ϵ on the first iteration, $\epsilon(1-\epsilon)$ on the second, and so on. The probability that all congregators have found each other within t iterations is:

$$\epsilon \sum_{i=0}^{t-1} (i+1)(1-\epsilon)^i = 1 - (1-\epsilon)^t. \tag{2}$$

We can then solve for t to ask how long it takes for a set of congregators to find each other with probability p.

$t \approx m^{\phi-1} |ln(1-p)|$.

Once all the agents have converged to the same congregation, they will remain together, since they are all receiving a positive payoff. The salient point in this example is that the time needed for agents to find each other increases exponentially with the number of agents in the system. We will use the time to converge as an indicator of the problem's hardness or complexity throughout the remainder of this paper.

We can use the above formulae to determine the corresponding CLRI parameters. We begin with c, the change rate, which is the probability that an agent will change an incorrect mapping. This will happen any time a congregator is not grouped with every other agent of its affinity group, so $c = 1-\epsilon$. Similarly, l, the learning rate, is the probability that an incorrect mapping will be changed to a correct one. This is simply c times the probability of picking the right locus, which is $(\frac{1}{m})$. r, the retention rate, is 1. When congregators find the right location, they quit moving and so there is no chance of "forgetting." Finally, we have volatility (v), which is the probability that the target function will change. This is the probability that the congregators are not all in the same location multiplied by the probability that the "right" locus changes on the next iteration: $v = (1-\epsilon)\frac{m-1}{m}$.

The above parameters indicate that as more agents are added, the problem becomes exponentially more complicated. In particular, v contains the terms $(1-\epsilon)$ and $\frac{m-1}{m}$. Since ϵ becomes exponentially small as the number of agents increases, volatility approaches $\frac{m-1}{m}$ in the limit. If m is at all large, this will be close to 1, meaning that every agent's actions have a very strong effect on the target functions of other agents. This is what we would intuitively expect, since the problem depends on all agents making the same decision.

4.2 Multiple Affinity Groups

We now generalize the problem to a case where there are g affinity groups of congregators, where each group is of size s, for a total of gs congregators. Again, there are a total of m loci (in the experiments, we will make the simplifying assumption that $m = g$). This means that there are m^{sg} different world states. Once again, each congregator wants to join a congregation which maximizes its payoff. In this experiment, we modify the payoff function slightly.

$$Payoff(a) = \begin{cases} \Sigma_{c \in C(a)} \frac{Distance(c)}{|C(a)|} - tax & \text{if } sum > tax \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where $C(a)$ is the congregation an agent is in (and $|C(a)|$ is the number of agents in this congregation) and tax is an "existence tax" that agents must pay every iteration (to pay for computational resources, etc.). $Distance$ is a function that determines how different (or how far apart) two affinity groups are. To ease analysis and remain consistent with the assumption discussed above that there be only one correct action in a given world state, we will let $Distance$ equal 1 if agents are in the same affinity group and 0 otherwise. In addition, to make non-optimal actions undesirable, we set tax equal to $\frac{s}{s+1}$. In short, this function assures that agents will only receive a positive payoff and therefore stay still in the case where they are in a congregation consisting solely of their own affinity group.

We are now able to determine the CLRI parameters. This problem is quite complicated, since there are such a large number of combinations of congregators. The first parameter is c, the change rate. A congregator will change an incorrect mapping any time it is not in the correct congregation, which, if all agents move randomly, will happen with probability $c = 1 - ((\frac{m-1}{m})^{(g-1)s})$. This is the probability that

[1]Derivations for this and following equations are left out for space reasons. They are available from the authors upon request.

a congregator is not in a congregation consisting solely of members of its affinity group.

The second parameter is the learning rate (l). This is simply the change rate c times the probability of randomly moving to the best available congregation, which is ($\frac{1}{m}$).

The third parameter is the retention rate (r). A congregator will retain a correct mapping if it is in the correct congregation and no congregators of another type arrive. In other words, $r = (\frac{m-1}{m})^{(g-1)s}$.

The fourth parameter is impact (i), which is the effect that agent b's decision has on agent a's target function. Agent b can affect agent a in two ways: if it is different and moves into agent a's congregation, or if it is the same and it moves out. The first case will happen with probability $\frac{1}{m}$. In the second case, a like agent will move out if there is a different agent in the congregation, which happens with probability $1 - \frac{m-1}{m}^{s(g-1)}$.

To make things concrete, assume that there are 3 affinity groups, each containing 3 agents. There are also 3 loci. This is the smallest example that demonstrates the complexity of the congregating problem. $c \approx 0.998$, $l \approx 0.332$, $r \approx 0.087$, the impact that like agents have on each other is $\frac{1}{3}$, and the impact that different agents have on each other is approximately 0.913.

Determining closed-form solutions for the time to convergence and the probability of convergence is more difficult in this scenario, since it is possible for a subset of congregators to settle early on and wait for others to find them. In other words, the probability of convergence is dependent on the entire past history. If we examine the learning rate and impact of agents from different affinity groups, we can see that the problem's complexity lies in the fact that both terms contain exponents for the number of agents: ($s(g-1)$). As in the previous problem, as the number of agents increases, the number of possible solutions explodes and l declines. In addition, each agent is moving around through loci trying to find the right combination , so v approaches 1 as every agent impacts the learning of other agents.

4.3 Experimentation

Since it is difficult to analytically determine the time needed for a system with multiple affinity groups to converge, we performed some simple experiments to develop an estimate. Using the same scenario as described above (3 loci, 3 affinity groups of three congregators each), we begin by correctly fixing 8 congregators and having 1 move. We then fix 7 and allow 2 to move, and so on. Agents move randomly between congregations until they settle on a locus where there are only members of their own affinity group (i.e. their payoff was 1). They then stay there unless "pushed out" by the arrival of a congregator from a different affinity group. Results (averaged over 50 trials) are shown in Figure 1.

As predicted in the analysis, the time needed for a system

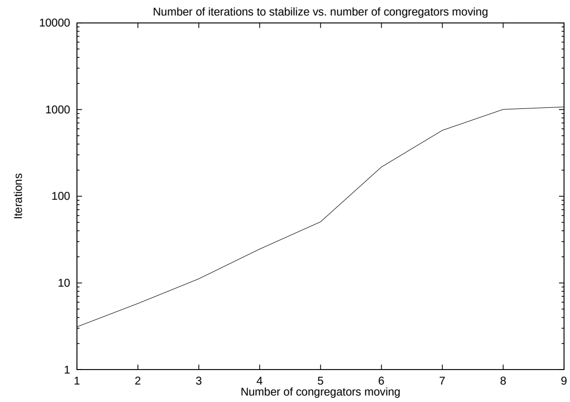

Figure 1. Iterations to stabilize as a function of the number of congregators moving.

of congregations to stabilize increases exponentially as the number of congregators increases. Even with this relatively small system, with a very simple payoff function, it quickly becomes very difficult for congregators to find each other.

5 Introducing labels

As we noted above, if agents are able to label the loci they are at and attract like members in that way, the congregating problem can potentially become much easier.

In this situation, the congregator's decision function becomes trivial; simply move to the congregation with the label corresponding to its affinity group. If no such label exists, move randomly. In this case, it is the labeler's problem that is the interesting one. Each labeler must choose a label which will attract an affinity group. The value of its decision will be influenced by the choices of the other labelers. Therefore, each labeler wishes to attract a distinct affinity group.

5.1 Analysis

In order to fit this problem within the CLRI framework, we make some simplifying assumptions. First, each labeler is neutral as to which affinity group it wants to attract. A labeler is happy whenever all congregators in its group are of the same type. Second, labelers have no memory regarding the payoffs from labels other than the current one. A labeler simply makes the decision as to whether to keep or discard the current label. We will also assume that each congregator agrees on the meaning of each label. If any congregator misinterprets the meaning of a label, they all will (in the same way).

We retain the notation used in the previous section. In addition, since there are m loci, there will be m labelers. Since g (the number of groups) is equal to m, there are m labels.

Within this framework, c (the change rate) is 1. If a labeler has an incorrect label, it will always change it. l, the learning rate, is the probability that the labeler will select an unused label, which is $(\frac{g-1}{g})^{m-1}$. r is also 1, by definition; if a labeler has a correct label, it will keep it. i is the impact that labeler i has on labeler j, which is the probability that i selects j's label, given that i was wrong and j was right. This is $((\frac{g-1}{g})^{m-1})\frac{1}{m}$.

Using the same parameters as above (3 labelers and loci, 3 affinity groups of 3 agents each), $l = \frac{4}{9}$ and the impact of i on j is $\frac{4}{27}$.

Once again, determining the time needed to converge is non-trivial, since one labeler's decision may aid another. In fact, what typically happens is that one labeler will attract an affinity group, who will then hold still and make the problem easier for other labelers. Therefore, to determine the time needed to converge with probability p, one must determine this probability for each possible history and then aggregate these probabilities appropriately.

The point to notice is that both learning rate and impact here are dominated by $(m-1)$ terms, meaning that the difficulty of the problem grows with the number of labelers and labels (or loci), rather than the number of congregations. In systems with a small number of labelers relative to congregators, labeling would seem to provide a substantial advantage.

5.2 Experimentation

Once again, since it was difficult to determine convergence time analytically, we decided to develop an experimental estimate. In this experiment, we began introducing labelers for each locus. A labeler could choose which affinity group it wanted to attract; its decision was independent of the other labelers. Congregators would automatically go to a group with the corresponding label (or choose randomly between offerings if more than one labeler offered the same label) or move randomly if the label for their group was not offered. Each congregator received a payoff according to the function described in section 4.2. The labeler's decision was analogous to the congregators' in the previous experiment: if the label chosen was successful (that is, it attracted a complete affinity group with no extra agents) it would continue to use that label. Otherwise, it would select a new label at random.

We began with a baseline experiment with no labelers. This is the same as allowing all congregators to move in the first experiment. We then added one labeler, then two, and then three. Results are shown in Figure 2.

There are a few things to notice here. The first is that the problem becomes trivial with three labelers. They settle on a solution almost immediately. Even with only two labelers, the problem becomes very easy; the congregators which don't have a label to attract them are able to settle on a common locus relatively quickly.

The slightly more subtle point is that adding one labeler produces the same effect as holding two congregators steady, and adding two labelers produces the same effect

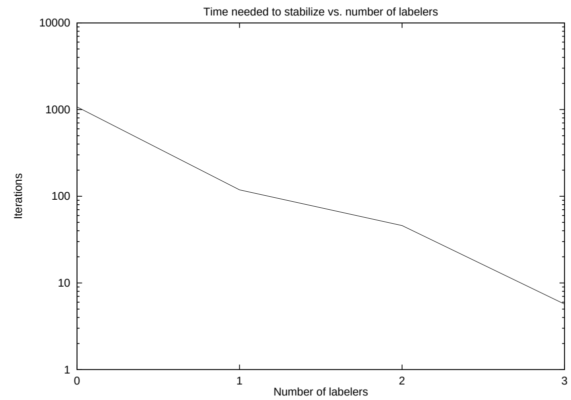

Figure 2. Iterations to stabilize as a function of the number of labelers.

as holding five congregators steady. The label produces a shared decision amongst all members of the corresponding affinity group; they will all move to the same congregation, and so in terms of convergence and impact they can be considered as one agent. This is a big savings; as was noted previously, a major difficulty with the congregating problem is the impact that one agent's decision has on another. If all agents in an affinity group make the same decision, then they will produce the same amount of impact on an agent of another affinity group as a single agent in a system without labelers. As more affinity groups are introduced, the potential benefits from introducing labelers will be even greater, since a system with labels will increase impact linearly (with each group added) rather than exponentially (with each group added *times* the number of agents in each group) in systems without labelers.

6 Discussion

The preceding analysis and experiments provide one piece of bad news along with one piece of good news. The bad news is that, as the number of agents which are trying to congregate increases, the problem becomes exponentially more difficult. In particular, an increase in the number of agents produces an increase in volatility, meaning that agents quickly begin to get in each other's way as they try to congregate. The good news is that introducing a basic coordinating mechanism such as labelers helps to reduce the complexity. When labelers are introduced, volatility grows with the number of affinity groups. There is still a point at which the problem becomes intractable, but labelers help to delay its onset. This idea of providing an external marker for congregators to find is similar to the idea of focal points [2]. Like labels, focal points help agents coordinate by providing an external mechanism for synchronizing behavior. Labelers can then be seen as agents which are responsible for the creation and deletion of focal points.

The introduction of labels transforms the problem from one in which each congregator must make a decision as to what congregation to join into one where each labeler must decide which label to offer. Labelers now must search through a space of labels in order to find the one which attracts the best congregation. This problem is typically easier to approach, for two reasons. First, the number of labelers is typically smaller than the number of congregators, and so coordinating decisions was easier. This was the point of the first experiment; the difficulty of coordinated decision-making grows exponentially with the number of agents making as decision. Second, there is often a structure to the space of labels that can be exploited during search. For example, labels may be related by a taxonomy or an ontology. Labelers which are aware of this structure can potentially use it to search the space of labels more efficiently.

We plan to extend this framework to other sorts of congregating domains, in particular information economies. As discussed in section 2.2, an information economy consists of producers (labelers) and consumers (congregators). The producers offer a set of goods and attract consumers who value those goods. A producer's problem is to select a set of goods (a label) which maximizes profit (attracts the optimal congregation). Given the results in this paper, it would appear that, since producers in an information economy must potentially interact with thousands of agents, they will have be able to efficiently use labeling information. In addition, if producers are competing with each other for consumers, their choice of labels can lead to the establishment of a niche market. By reasoning effectively about which label to choose, and therefore which niche to target, producers can increase their profits and avoid engaging in costly price wars. There are, of course, some significant differences between the affinity group domain and the information economy domain: for example, outcomes are not all-or-nothing, and the set of available labels is much richer. Extensions to contend with this are discussed in the following section.

7 Extensions

One weakness in applying CLRI to congregating is that it only considers pairwise impact. In congregating, the impact of one agent on another may also depend upon the presence of a group of agents. We would like to extend the framework to capture the impact that groups of agents can have on each other.

Also, the experiments described in this paper use congregators which move randomly. This makes analysis tractable, but it is not particularly realistic. Therefore, we plan to incorporate congregators which actually learn based on past history into our model. Also, both congregators and labelers should make decisions based not only on their myopic reward, but rather upon a consideration of their discounted long-term reward.

The basic CLRI model has no notion of utility-based actions; there is one correct action and all others are incorrect. In congregating problems, a finer granularity is needed; we would like to talk about one congregation being preferred to another, and by how much. This would allow us to consider the question of when a "good enough" set of congregations has been found; should congregators stay in a not-quite-optimal system of congregations, or expend more energy learning in the hope that the optimal system will produce more reward?

Finally, in domains such as information economies which have a large number of congregators and labels, a labeler may be able to exploit relations between labels, such as an ontology or a taxonomy, so as to eliminate certain labels from consideration and settle more quickly on a label which attracts a desirable congregation.

References

[1] C. H. Brooks and E. H. Durfee. Congregation formation in information economies. In *Proceedings of the AAAI-99 Workshop on AI in Electronic Commerce*, 1999.

[2] M. Fenster, S. Kraus, and J. Rosenschein. Coordination without communication: An experimental validation of focal point techniques. In *Proceedings of the First International Conference on Multi-Agent Systems*, pages 102–116, San Francisco, California, 1995.

[3] P. J. Gmytrasiewicz and E. H. Durfee. Toward a theory of honesty and trust among communicating autonomous agents. *Group Decision and Negotiation (Special Issue on Distributed Artificial Intelligence)*, 2:237–258, 1993.

[4] J. Hu and M. Wellman. Multiagent reinforcement learning: Theoretical framework and an algorithm. In *Proceedings of the Fifteenth International Conference on Machine Learning (ICML-98)*. Morgan Kaufmann, 1998.

[5] T. M. Mitchell. *Machine Learning*. WCB/McGraw-Hill, Boston, Massachusetts, 1997.

[6] J. S. Rosenschein and G. Zlotkin. *Rules of Encounter: designing conventions for automated negotiation among compters*. MIT Press, Cambridge, MA, 1994.

[7] S. Sen, M. Sekaran, and J. Hale. Learning to coordinate without sharing information. In *Proceedings of the National Conference on Artificial Intelligence*, pages 426–431, 1994.

[8] O. Shehory and S. Kraus. Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101:165–200, 1998.

[9] K. Sycara, K. Decker, A. Pannu, M. Williamson, and D. Zeng. Distributed intelligent agents. *IEEE Expert*, 11(6):36–46, December 1996.

[10] M. Tambe. Toward flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.

[11] J. M. Vidal and E. H. Durfee. The moving target function problem in multi-agent learning. In *Proceedings of the Third Annual Conference on Multi-Agent Systems*, Paris, France, 1998.