Effect of referrals on convergence to satisficing distributions

Teddy Candale
Math & Computer Sciences Department
University of Tulsa
teddy-candale@utulsa.edu

Sandip Sen
Math & Computer Sciences Department
University of Tulsa
sandip@utulsa.edu

ABSTRACT

We investigate a framework where agents locate high-quality service providers by using referrals from peer agents. The performance of providers is measured by the satisfaction obtained by agents from using their services. Provider performance depends upon its intrinsic capability and upon its current load. We present an algorithm for selecting a service provider for a given task which includes mechanisms for deciding when and who to ask for a referral. This mechanism requires learning, over interactions, both the performance levels of different service providers, as well as the quality of referrals provided by other agents. We use a satisficing rather than an optimizing framework, where agents are content to receive service quality above a threshold. Agents have to learn the quality of others' referrals and the quality of providers to find satisficing providers. We compare the effectiveness of referral systems with or without deception with systems without referrals. We identify zones, based on an observed entropy metric, where using referrals is helpful in promoting fast convergence to satisficing distributions.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence[Coherence and coordination, Multiagent systems, Intelligent agents]; I.2.6 [**Artificial Intelligence**]: Learning—*Knowledge acquisition*

General Terms

Performance, Experimentation

Keywords

load balancing, referral system, satisficing distribution, simulated environment

1. INTRODUCTION

Location of high-quality services and load balancing are challenging problems in a large number of multiagent systems. Although both involve location of high-quality providers, the nature of resources is different in the two cases. Service location addresses environments where agents need to identify resources or peers with the required expertise to answer queries [7, 11, 13, 19, 20]. An often-used approach is for agents to ask other agents for referrals to services and resources. In a number of service location domains, the cost of referring is generally assumed to be negligible for a referrer and the increased load on the corresponding service is not assumed to lead to a decrease in its performance. While recommending a web-search engine, for example, a higher rate of usage, for the most part, does not produce a decrease in response time, and popularity can be used as an indirect measure of quality. In other domains, referring others to one's preferred service can increase the latter's revenue and in turn even increase the performance of the service due to increased or improved resources financed from the additional revenue.

On the contrary, problems solved using load balancing present environments comprising fixed number of resources whose performances are directly related to their workloads. An Agent's choices of a resource affect both its own utility and that of other agents currently using the chosen resource. Besides, in this context, referrers, or agents working in a coalition, may incur a non-negligible cost due to an increase in load on their preferred resources [16].

We investigate a model presenting similarities with the latter described kind of environments. The performance of a provider depends both on its intrinsic characteristics and the current workload it is handling. There exists a limited number of service providers in the environments requiring agents to at least implicitly coordinate their selection of service providers. Myopic, self-interested behavior can lead to poor performance for the individual and system-wide instability. There is thus a need for non-myopic mechanisms to promote performance and stability of such decentralized agent systems. In particular, deriving protocols and strategies that lead to equilibrium states where all agents in the community are satisfied is a challenging and significant research problem.

Referrals from other agents can help agents find more satisfying service providers. But such referrals may cost the referring agent since the load of the referred provider may increase, with corresponding performance deterioration. This is particularly true if referral chains exist, i.e., if an agent can refer providers it located through referrals to other agents. Another problem with the use of referrals is that the recipient of a referral can be misled by deceptive referrers who deliberately provide false information to eliminate competition for its preferred providers. While referral systems have been widely studied both in theory and in practical applications [5] the negative side-effects of referrals have not received adequate treatment. We seek to analyze the benefits and disadvantages of referrals in domains where the cost of referrals is uncertain. The goal is to identify situations where an agent should or should not use referrals.

While ideally speaking agents may aspire for optimal satisfaction levels from service providers selected for performing an assigned task, dynamic, partially known, and open environments can render the realization of this ideal behavior improbable. Possible sources of inefficiencies include noisy, variable feedback about provider performance as the environment is at best partially observable which implies all factors affecting performance are not directly observable. In a dynamic environment the expected performance of a provider as referred by another agent may have changed based on current load and is not necessarily an indication of the trustworthiness of the referring agent. Besides, an agent is unable to accurately assess the impact of its own decisions, including choice of service providers and making referrals, on its environment.

As such it might not be feasible to seek strategies for optimizing performance. Rather, we posit that agents should concentrate on finding service providers that provide a quality of service which exceeds an acceptable performance threshold. This formulation is consistent with Simon and others view of bounded rationality of decision makers within the context of complex organizations [4, 14]. Approaches from game theory also use the notion of *aspiration levels* to stabilize systems and reach equilibrium [2, 15, 17, 18, 3].

In this paper, we provide an approach for trading referrals using which agents can locate high-quality service providers. Our goal is to develop strategies by which a system of autonomous agents can quickly reach stable configurations where all agents are satisfied with the choice of their current service providers. Our proposal involves learning to rate referrers and use such ratings to adjust future referrals to identify effective service providers. We formally demonstrate the convergence to satisfactory service provider selections for the entire group of peer agents. Through our analysis we explain when an agent should or should not choose service providers based on referrals from other agents.

The paper is structured as follows. Section 2 introduces the environment and algorithms used by our agents. Section 3 is an analysis of criteria governing the speed of convergence while our experimental results are presented in Section 4. Section 5 briefly states some related works. Section 6 concludes the paper and introduces our future work.

2. FRAMEWORK

2.1 Environment

We present an environment where agents share a set of service providers to perform daily tasks. When an agent chooses a provider, it puts an associated load on the provider. They cannot directly observe or measure the intrinsic quality of these providers. After choosing a provider on a given day, they can observe their performance, which is dependent on the load on that provider on that day, at the end of the day. Agents are self-interested and are only interested in their own satisfactions.

We now introduce a formal definition of our environment. Let $\mathcal{E} = <\mathcal{A}, \mathcal{R}, \mathit{perf}, L, S, \Gamma>$ where:

- $\mathcal{A} = \{a_k\}_{k=1..K}$ is the set of agents
- $\mathcal{R} = \{r_n\}_{n=1..N}$ is the set of providers.
- $f: \mathcal{R} \times \mathbb{R}_+ \rightarrow [0, 1]$, provides the intrinsic performance of a provider given a load.
- $L : \mathcal{A} \rightarrow \mathbb{R}_+$, is the load function for the agents.
- $S: \mathcal{A} \times [0, 1] \rightarrow [0, 1]$, is the satisfaction function of agents.
- $\Gamma = \{\gamma_1, \ldots, \gamma_K\}$, is the set of satisfaction thresholds, representing aspiration levels, of agents.

Each day d, agent a_k is assigned a load $L(a_k)$. At the outset, each agent knows the set of providers that can process its task without the knowledge of their intrinsic capabilities represented by their performance function, $f(r_n, \cdot)$, for provider r_n. The agents are also unaware of the current load on the providers. We assume that the performance of a provider on a particular day depends on the total workload on that day: agents which use a provider the same day will obtain identical performances. Any two agents, however, may have different aspiration levels for the same quality of performance. We represent the satisfaction of an agent a_k by a *subjective* function $S(a_k, \cdot)$ which models the satisfaction obtained based on the quality of service provided by the chosen provider. To be more precise, if $\mathcal{A}_n^d$ is the set of agents using the provider r_n at day d then the provider's performance after processing all these orders is $\mathit{perf} = f\Big(r_n, \sum_{a \in \mathcal{A}_n^d} L(a)\Big)$. perf is the service quality received at the end of the day d by every agents in $\mathcal{A}_n^d$. An agent $a_k \in \mathcal{A}_n^d$ will evaluate the performance of r_n by the value $s = S(a_k, \mathit{perf})$. s is the satisfaction obtained by agent a_k using provider r_n. This agent will be satisfied if this performance s is above the threshold γ_k. This model allows two agents to have different satisfaction levels for the same quality of service received.

Agents are interested in obtaining satisfactions above a threshold rather than maximizing it. Our aim is to design interaction protocols and behaviors that allow all agents to find satisfying providers. The following definitions formalize this idea.

Definition 1 (Distribution of agents over providers). *We call* **distribution of agents over providers** *the set* $D = \{\mathcal{A}_n\}_{n=1..N}$ *such that:*

- $\mathcal{A}_n \subseteq \mathcal{A}$ *($\mathcal{A}_n$ may be empty)*
- $\bigcup_{n=1}^{N} \mathcal{A}_n = \mathcal{A}$
- $n_1 \neq n_2 \Longrightarrow \mathcal{A}_{n_1} \bigcap \mathcal{A}_{n_2} = \emptyset$

The set of distributions is denoted by $\mathcal{D}$.

DEFINITION 2 (Γ-ACCEPTABLE DISTRIBUTION). *A distribution $D = \{\mathcal{A}_n\}_{n=1..N}$ is said to be* **Γ-acceptable distribution** *iff*

$$\forall n,\, a_k \in \mathcal{A}_n \Longrightarrow S\left(a_k,\, f\left(r_n,\, \sum_{a\in\mathcal{A}_n} L(a)\right)\right) \geq \gamma_k$$

. The set of Γ-acceptable distributions is denoted by $\mathcal{D}_\Gamma$.

The concept of distribution represents how agents distribute themselves over the providers. For example, $\mathcal{A}_n$ is the set of agents who use the provider r_n. A Γ-acceptable distribution is a distribution where every agents receive a satisfaction above their own satisfaction threshold. A Γ-acceptable distribution is expected to be a stable distribution since no agent will have the incentive to change their choice of provider for the next day as each of them is satisfied. Consequently, it is an equilibrium concept and our goal is to enable agents to reach such distributions.

2.2 Referral based provider selection

We present alternative strategies for selecting service providers. We assume that apart from the feedback received in the form of performance of chosen providers, agents can receive referrals from other agents stating their satisfaction with different providers.

The quality of a provider is evaluated by calculating the mean of satisfactions an agent has obtained from a provider. This value is the satisfaction an agent may expect to get by using the provider again.

We believe that improving system stability is necessary to ensure faster convergence and also helps agents to obtain more accurate information about their environment. Therefore, we use a "move only when you think you can do better"-principle for agent strategies. We evaluate three kinds of agents: agents who find providers on their own without using information from other agents (*No Referral* or *NR*), agents who use referral to locate providers and are trustful of the referrals received(*Referral (Truthful)* or *RT*), agents who use referral to locate providers but always deceive while giving referrals (*Referral (Deceptive)* or *RD*).

NR: Agent a_k finds the provider by using only its own experience and without referrals. This agent never moves from a provider if it is satisfied. Otherwise, either it explores, i.e., chooses a provider randomly, with a probability α or with a probability $1-\alpha$, it makes a decision in the following way. It estimates the expected satisfaction, es, it can get from every provider. It picks a provider among those for which $es > \gamma_k^-$, a constant less than γ_k. If no such provider exists, a_k does not move.

RT: Agents may help each other by giving referrals. A referral includes a provider name and an estimation of its quality. Before asking for referral, the asking agent determines the set $\mathcal{A}_h$ of agents whose expected quality of referral is greater than γ_k^-. It picks one of them, a_h, with probability proportional to q_h, the expected quality of referral from that agent. The referral will be accepted if the provided estimation is greater than γ_k^-. If not, the asking agent will choose another referrer in $\mathcal{A}_h$. The asking agent trusts the information given by the referrer and does not try to correct the estimation. Agents are assumed to be helpful: they refer only the best providers in their estimate. They are also assumed to be truthful, providing the true estimation of the provider performance. If a referring agent knows of no good provider, one is given randomly as referral but with estimation 0. Agent a_k chooses a provider by using the same protocol as defined in NR except that if it does not find a satisfying provider using its own information, it asks for referral from an agent chosen either (a) randomly with probability α_{ask} (exploration) or (b) using estimation of other agents usefulness with probability $1-\alpha_{ask}$.

RD: This case is identical to the previous one except the fact that agents always deceive. They refer the same provider as *RT*, but alter the true estimation, es. More precisely, if es is greater than their satisfaction threshold then they report an underestimation, otherwise an overestimation.

3. CHARACTERIZING SYSTEM CONVERGENCE

In an environment where agents are autonomous and self-interested, they are expected to choose actions which appear to be the best ones to them. Such greedy, myopic individual actions can lead to conflicts that affect system performance and, in turn, reduces individual satisfaction. The present case of agents choosing satisfactory service providers is a clear example of such situations as agent satisfactions can oscillate with changing loads on satisfactory providers. Such systems tend to oscillate between good and bad distributions produced by slow convergence and variable agents satisfactions. We believe this phenomena is appropriately modeled by analyzing the effect of agent *inertia* on the system.

3.1 Influence of inertia

We believe that oscillations in our environment will happen if a system at a distribution close to a Γ-acceptable distribution will have the tendency to evolve to a worse distribution and vice versa. Here we assume that the total load to be applied by all agents in the system is approximately equal to the total capacity of all service providers to produce satisfactory performance for all agents if they are properly distributed. Intuitively, a distribution where almost everyone is satisfied contains very few under-used or over-used providers and the rest are occupied by the right number of agents. Those under-used providers $\mathcal{R}_u$ are very attractive. Consequently, agents will be inclined to move to them, which leads the system to a distribution where providers in $\mathcal{R}_u$ will be overcrowded. We believe this key, problematic effect can be mitigated by increasing the inertia in the system, where inertia is an inverse function of the number of agents moving at any given time.

We will now formalize this analysis. In the remainder of this paper, we assume our agents are identical in the sense they have the same satisfaction functions ($S(a_k,\cdot) = S,\ \forall k$), the same satisfaction thresholds ($\gamma_k = \gamma,\ \forall k$) and the same daily load ($L(a_k) = L$). Consequently, we can define the capacity the providers as the maximum number of agents they satisfy at the same time. A Γ-acceptable distribution is then a distribution where every providers is used by a number of agents less than or equal to their capacities. We measure the goodness of a distribution as the number of agents who should move to reach a perfectly coordinated distribution. We now give a formal definition of the measure we call entropy.

DEFINITION 3 (ENTROPY). *Given an environment where provider r_n has its capacity equal to C_n and agents are iden-*

tical, we represent entropy of a distribution D by

$$\mathcal{E}(D) = \sum_{n=1}^{N} \max\Big(0,\ |\mathcal{A}_n| - C_n\Big).$$

We can see that each Γ-acceptable distribution has an entropy equal to 0. The lower the entropy the better the distribution. Proposition 1 shows how the number of moving agents influence the entropy of the system.

PROPOSITION 1. *If K^d_{move} is the number of agents which can move at date d then*

$$\mathcal{E}(D^{d+1}) \in [\max\big(0,\ \mathcal{E}(D^d) - K^d_{move}\big),\ \mathcal{E}(D^d) + K^d_{move}]$$

Proposition 1 shows reducing K_{move} has a beneficial effect of the entropy of the system. Indeed, it determines the size of the range in which the next entropy of the system belongs. Consequently, we may expect oscillations of the system entropy with low amplitude or no oscillation at all for small values of K^d_{move} and high amplitude oscillation for high values of K_{move}. Our agents are autonomous in their decision-making, no central decision can be taken. Therefore, we claim that a high inertia implies low K_{move} in average. Nevertheless, K_{move} is also influenced by the number of agents which are inclined to move (K_{wm}). In other words, agents who are not satisfied by their current providers will seek to switch. K_{move} is positively correlated to K_{wm}. Proposition 2 expresses that when the number of agents is fixed, the size of the range of K_{wm} is inversely correlated to the number of providers N. Intuitively, the arrival of a new agent in provider r_n with current load equal to its capacity C_n makes $C_n + 1$ agents envisage the possibility to go to another provider.

PROPOSITION 2. *Consider an environment where providers have a capacity of K/N. Then the number of agents inclined to move is in the range*

$$K_{wm} \in \Big[S(D) + \frac{K}{N},\ S(D)\cdot(\frac{K}{N}+1)\Big].$$

Initial experiments showed the existence of an upper bound K^*_{move} such that the system is stable if $K_{move} \leq K^*_{move}$. From this observation and Proposition 2 we conclude that in an environment where providers are identical, and have capacity K/N, there exists a lower bound N^* of the number of providers such that stable systems with satisfactory distributions are obtained for $N \geq N^*$. In other words, we may expect two zones where the behaviors of our algorithms will be fundamentally different. In Zone I (defined by $N < N^*$), we may expect poor performance of the system: low speed of convergence and variable received satisfaction for agents. In Zone II (defined by $N \geq N^*$), we expect to obtain excellent or acceptable performance: reasonable or fast speed of convergence and consistent satisfactory provider performance.

3.2 Exploration

The previous section establishes that the system performs better when few agents move at any given time. As agents move less frequently, it is more likely that their decisions are based on more accurate information. Correspondingly, more informed decisions will expedite system convergence to satisfactory distributions. But such convergence also requires learning about provider and referral qualities. Consequently, some systematic exploration of providers and referrers appear to be necessary. On the other hand, such exploration decreases the inertia of the system and can impact convergence rate. An environment where agents explore too much will produce system instability where agents will hardly have representative estimations of provider performances since loads vary significantly. As a result, agents will not receive consistent satisfying provider performance, giving them more incentive to move. In this context, referral systems seem to be appealing since agents may substitute their exploration with others' experience.

A problematic issue that is often ignored in referral systems is the long-term cost which can be incurred by referring a service whose quality is inversely proportional to usage. We can assess this referral cost locally by the loss of satisfaction for individual agents and globally as the time needed to reached a Γ-acceptable distribution for the entire community of agents. In the case of benevolent and truthful referrers, helping agents are directly affected since their referrals increase the load of their favorite service providers. It may be also harmful for the other agents who use referred providers and then for the system since the number of agents who are predisposed to move may increase, decreasing system inertia, and increasing entropy and convergence time.

Sen et al's [10] explanation of this phenomenon is that by asking for referral, the amount of information available to an agent increases, which, in turn, can be detrimental to system stability in such domains.

4. EXPERIMENTAL RESULTS

We present results from experiments designed to evaluate the relative merits of using referrals to choose service providers. All experiments were run with 200 identical agents with a satisfaction threshold equal to 0.7 ($\forall k,\ \gamma_k = 0.7$), with $\forall k, S(a_k,\ x) = \frac{1}{1+0.7\,x}$. We choose $f(r_n,\ L) = (\delta_n \cdot L+1)^3 - 1$ as provider performance functions. The parameter δ_n is tuned to obtain providers with desired capacities. Each day an agent is assigned a task whose load is 1 ($\forall k,\ L(a_k) = 1$).

4.1 Uniform provider capacities

In this subsection, we study the influence of the number of providers available and consequently the total provider capacity. All providers have the same capacity equal to K/N such that a Γ-acceptable distribution is a distribution where each provider is used by exactly K/N agents. Table 1 presents the average number of iterations needed to reach the convergence over 50 runs. We limited the number of iterations in each run to 5000. In the table `na` stands for "not available" and corresponds to runs where convergence was not reached within 5000 iterations.

We highlight the following observations:

1. Results confirm our prediction of the existence of the expected lower bound, N^* for effective system performance. For any agent type, convergence speed is optimal when $N = N^*$. The estimates for the different agents are: $N^*(NR) = 40$, $N^*(RT) = 100$, $N^*(RD) = 100$.

2. For any agent type, performance in Zone I, i.e., for $N < N^*$ is worse compared to performance in Zone II,

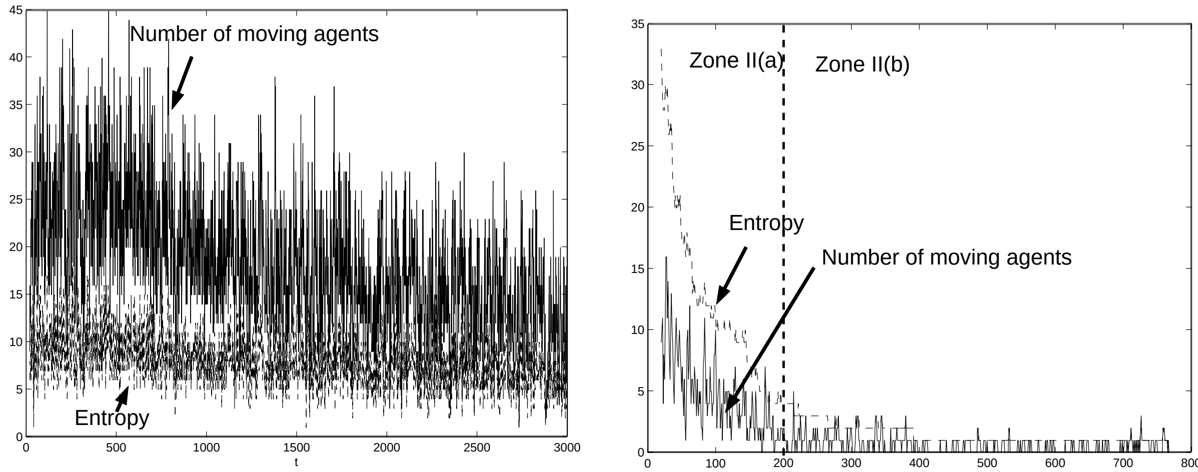

(a) Evolution of K^t_{move} and $\mathcal{E}(D^d)$ in Zone I (10 providers, NR)

(b) Evolution of K^d_{move} and $\mathcal{E}(D^d)$ in Zone II (200 providers, NR)

Figure 1: Examples of evolution of K^d_{move} and $\mathcal{E}(D^d)$ in Zones I (left) and II (right).

N	NR	RT	RD
200	2206.529	1578.863	1867.706
100	624.588	454.510	558.078
40	167.647	3002.059	na
20	3624.647	na	na
10	3879.471	na	na

Table 1: Average iterations to convergence.

i.e., for $N \geq N^*$.

3. When $N \geq N^*$ for all strategies, i.e., in the range $N \geq 100$: *RT* converges faster than *RD*, which converges faster than *NR*.

4. *NR* is more robust than other algorithms as it produces convergence for a much larger range of environments, e.g., only *NR* leads to convergence within the iteration limit for $N \leq 20$.

An important observation that builds on the above points is that: *While the use of referrals from truthful agents can speed up system convergence to satisfactory distributions, such knowledge sharing can also increase system entropy and slow down convergence with a relatively small number of providers in the environment.*

To obtain a deeper understanding on the nature of convergence with and without referrals and for different number of providers, we studied the K^d_{move} and $\mathcal{E}(D^d)$ metrics over the course of different runs. On inspecting these metrics for different system configurations we find that even though the system convergence systems worsens when we move away from $N = N^*$, there is an interesting, clear difference between runs corresponding to Zones I and II.

Figure 1 shows the evolution of the entropy and the number of simultaneous moving agents with 10 providers (corresponding to Zone I) and with 200 providers (corresponding to Zone II) for *NR*. We include only the graph for *NR* as graphs for *RT* and *RD* are similar. In Zone I, the entropy, $\mathcal{E}(D^d)$, keeps on oscillating while K^d_{move} remains higher than it. Consequently, the system cannot converge monotonically as more agents than required move at the same time. This contrasts with the behavior of the system in Zone II: K^d_{move} remains mostly below $\mathcal{E}(D^d)$ preventing system instability and promoting convergence. We can further differentiate Zone II runs into two parts: (i) in Zone II(a) the system moves consistently toward almost coordinated distributions, (ii) in Zone II(b) the system remains in almost coordinated distributions and a small number of unsatisfied agents keep moving in search of a satisficing provider. Rustogi & Singh experimentally show that convergence in similar systems, but without referrals, can be improved when tolerating imprecision [8]. Assuming tolerating imprecision is equivalent to claiming convergence when a distribution with a small non-zero entropy is obtained, and given the fact that the system quickly reaches Zone II(b), 'convergence' can be significantly expedited by 'tolerating imprecision'.

Though iterations to convergence for *NR* (see Table 1) appear roughly equivalent in Zones I and II, e.g., when $N = 200$ and $N = 10$ respectively, the satisfaction levels of individual agents are fundamentally different. For $N = 200$, corresponding to Zone II, almost the entire community is in a satisficing state, i.e., entropy is low for a significant portion of the run. Exploration of few agents needing to improve satisfaction does not destabilize the system. On the contrary, for $N = 10$, corresponding to Zone I, most agents are unsatisfied and their explorations of different service providers, to improve their own satisfaction levels, hurt the entire society, thus delaying convergence.

The reason of the better performances of *RT* when $N^* \leq N$ is because, in Zone II(b), the exploration requirement of the few unsatisfied agents looking for better providers is partially reduced to by referrals. As a result, unsatisfied agents need less time to find satisfactory providers.

It appears that *RT* and *RD* face unique convergence for small number of providers. This is actually a problem of

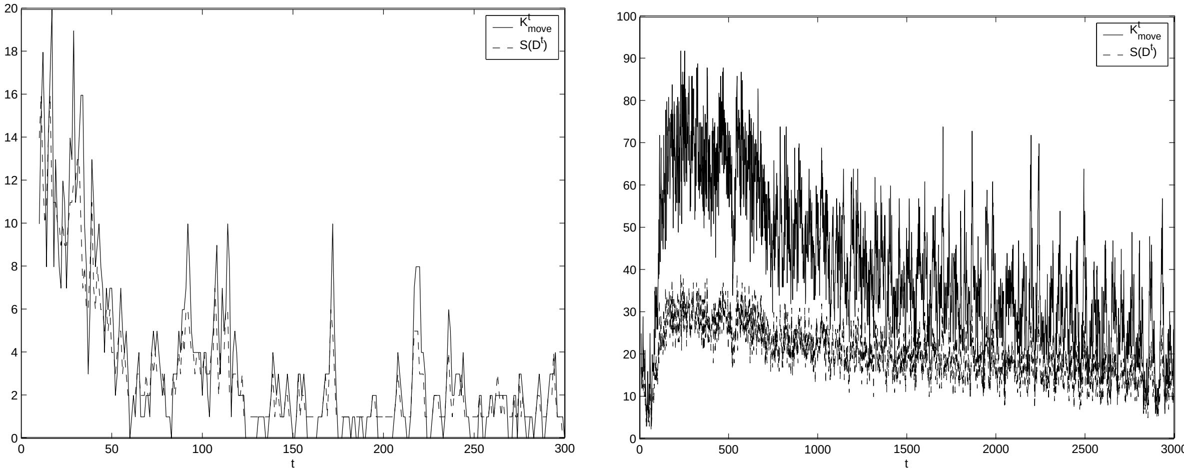

(a) 40 providers (*NR*)

(b) 40 providers (*RT*)

Figure 2: K^d_{move} and $S(D^d)$ for NR (left) and RT (right) for N=40.

scale-up with the number of agents, which is isomorphic to decreasing the number of providers keeping the number of agents constant. As illustrated in Figures 2(b) and 2(a), K_{move} is too high for *RT* whereas *NR* manages to limit that number. *RT* cannot stabilize the system, and this prevents the system from reaching convergence. We had predicted that in environments with few providers, agents are more inclined to move. Two factors are responsible for less inertia of *RT*: the use of referrals and the amount of information available to a particular agent. Recall our agents do not move if they do not think they will improve their satisfaction. The use of referrals augments the probability to move because of the agents' trust in the referrer. Besides, by visiting the referred provider, an agent increases the number of providers it knows, which makes it more inclined to move in future iterations. Consequently, in systems with referrals, agents have more incentive to move leading to more instability.

4.2 Non-uniform provider capacities

We now present, in Table 2, results from experiments when providers have different capacities. The environment contains a large number of providers with low capacities (equal to 2) and few providers with high capacities (equal to 20). We observe that

• As before, convergence of *NR* is slower than that of *RD* which is slower than that of *RT*.

• When the number of providers with high capacity is increased, all strategies perform better.

The last observation can be explained by the fact that, in the beginning, agents have the tendency to spread out and all providers have almost the same load. Agents need time to realize that some of them need to move to providers with high capacity. This exploration time is reduced when more high-quality providers are present. Besides, with referral, agents take less time to find high-capacity providers (shorter size Zone II(b)). Negative side-effect of referrals is less in this situation as agents using high-capacity providers can refer them with less risk of their load increasing to the extent that the referrers satisfaction will drop below their satisfaction threshold. However, *RD* performs worse than *RT* since agents of the former type mislead each other by providing wrong estimations of provider quality. Hence, the resource discovery process (Zone II(b)) is not accelerated compared to environments where only *NR* is present.

# providers with capacity:		*NR*	*RT*	*RD*
20	2			
1	90	1059.608	834.137	1028.177
5	50	300.235	229.882	271.725

Table 2: Average number of iterations to reach the convergence (200 agents).

5. RELATED WORK

Economic approaches often focus on the use of negotiation to determine acceptable resource trades [1, 6, 9]. Agents possess sets of resources that they can trade with other agents. An agent is assumed to use resources it owns. This contrast with our setting where agents can use any resource in the system. Negotiation outcomes are evaluated using two concepts: *utilitarianism* and *egalitarism*. The utilitarian concept consists in the maximization of the sum of utilities of all agents present in the society. It is regarded as unfair since it often produces high variability in individual utilities. In the optimal allocation, an agent may have a high utility while another one has a very poor one. The egalitarian concept tries to maximize the utility of the agent with the lowest utility. Our concept of Γ-acceptable distribution can be related to this latter concept. A state is regarded as acceptable if every agent is ensured to receive a minimum acceptable satisfaction.

Referral systems have recently received increasing attention among multiagent researchers. In [19], Yu and Singh study a referral system when an agent helps a human user find relevant expertise and protect him/her from too many irrelevant requests. The agent has to learn both its associated user's and others' expertise. Interactions and chains of referrals are used to update other's expertise. When modifying them the sociability of peer agents is used to favor reciprocative agents. Sen & Sajja have studied the use of referrals to locate service providers when an agent first enters a new community with no prior knowledge of the quality of service providers or the reliability of the referrers [11, 12]. In those previous works, peers have a short term cost of processing the referral request, which can be negligible in most domains. In our setting, referrals have a long term cost as the asking agents may use the referred provider in the future and also refer it to others and hence possibly reduce the performance of that provider. Thus, chains of referrals may result in a consistent decrease in the helping agent's received satisfactions.

Coordination is a key issue in multiagent systems. Sen, Arora, and Roychowdhury [10] show that information can negatively impact agent coordination to find balanced distribution among resources. They allow agents to move to providers only in the neighborhood (called window) of the one they are currently using. They achieve perfect coordination faster when the window size decreases. Using a probabilistic analysis, they demonstrate agents are much more inclined to move from an overcrowded provider when the window size is high and the opposite when the size is low. Besides, an under-used provider is likely to become over-used at the next time step when the size window is high. They conclude that too much information available to agents lead to oscillating provider loads. This leads to variable provider performances and low speed of convergence. Rustogi & Singh [8] study the influence of inertia for system convergence in the same domain. Inertia is an agent's reluctance to move even when it believes that it can do better with a different resource choice. They proved that high inertia speeds up convergence when knowledge increase but low inertia perform better with little knowledge.

We choose the domain used by both of these papers to study the merits and de-merits of providing referrals in systems where referrals have uncertain, long-term costs. We believe that a more comprehensive understanding of system behavior can be obtained by studying the number of simultaneously moving agents, K_{move}. Our approach provides more detailed characterization of the system, but is consistent with general conclusions from Rustogi & Singh [8] study as inertia can be used as a parameter controlling K_{move}. Rustogi & Singh [8] also claim that performance is remarkably improved when accepting imprecision. Our analysis in the last section shows that this happens because, in Zone II, low entropy is reached quickly, whereas a much larger number of iterations is needed for perfect coordination. By tolerating imprecision or declaring convergence when only a small number of agents are unsatisfied, i.e., when entropy is low, convergence is significantly accelerated.

6. CONCLUSION AND FUTURE WORK

We have investigated the benefit of referrals to locate service providers where referrals may have long-term cost to the referrers by increasing the load and thereby decreasing performance of their preferred service providers. But referrals also help locate providers that can improve satisfaction of the recipient agents. The research question is when and how does this cost-benefit tradeoff improve or harm convergence rates to satisfactory distributions. Interestingly, systems without referrals appear to be more robust in the sense they have satisfactory or reasonable performance even for extremely small number of providers, i.e., for more challenging environments. Referrals, however, do facilitate convergence when there are a significant number of providers. Deceitful referrers, unfortunately, can slow convergence. This can be remedied if recipients learn the truthfulness of the referrers.

We present the relation between system entropy and number of concurrently moving agents as a key characteristic underlying convergence rates. The latter number is governed both by agent inertia and by the number of providers in environments with uniform providers. We identify the existence of two zones where the agent satisfaction levels are noticeably distinct. In one zone, with large number of providers, convergence speed is acceptable and consistent obtained satisfactions to the entire society even before convergence. In the other zone, with few providers, slow convergence is observed with highly variable individual satisfactions.

A key observation from our analysis is that agent exploration should be limited to improve convergence rate. We plan to design an algorithm allowing agents to adjust their exploration rate given their perception of the stability of the system. Another area of study will be to make agents learn faster by not only evaluating the satisfaction they can get from providers but also the intrinsic capabilities of the provider.

Acknowledgments: This work has been supported in part by an NSF award IIS-0209208.

7. REFERENCES

[1] U. Endriss and N. Maudet. Welfare engineering in multiagent systems. In A. Omicini, P. Petta, and J. Pitt, editors, *Engineering Societies in the Agents World IV*, volume 3071 of *LNAI*, pages 93–106. Springer-Verlag, 2004. Postproceedings of ESAW-2003.

[2] G. E. Kersten. NEGO - group decision support system. *Information and Management*, 8:237–246, 1985.

[3] V. Lotfi, T. J. Stewart, and S. Zionts. An aspiration-level interactive model for multiple criteria decision making. *Comput. Oper. Res.*, 19(7):671–681, 1992.

[4] J. G. March and H. A. Simon. *Organizations*. John Wiley and Sons, New York, 1958.

[5] D. W. McDonald. Recommending collaboration with social networks: a comparative evaluation. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 593–600, New York, NY, USA, 2003. ACM Press.

[6] H. Moulin. Axioms of cooperative decision making. Cambridge University Press, 1988.

[7] M. H. MP Singh. *Service-oriented computing: semantics, processes, agents*. John Wiley & Sons, 2004.

[8] S. K. Rustogi and M. P. Singh. Be patient and tolerate imprecision: How autonomous agents can coordinate effectively. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 512–517, 1999.

[9] A. Sen. Collective choice and social welfare. Holden

Day, 1973.

[10] S. Sen, N. Arora, and S. Roychowdhury. Effects of local information on group behavior. In *Proceedings of the Second International Conference on Multiagent Systems*, pages 322–329, Menlo Park, CA, 1996. AAAI Press.

[11] S. Sen and N. Sajja. Robustness of reputation-based trust: boolean case. In *AAMAS '02: Proceedings of the first international joint conference on Autonomous agents and multiagent systems*, pages 288–293, New York, NY, USA, 2002. ACM Press.

[12] S. Sen and N. Sajja. Selecting service providers based on reputation. In *Proceedings of the AAAI-2002 Workshop on Multi-Agent Modeling and Simulation of Economic Systems*, 2002.

[13] S. Sen and N. Sajja. Selecting service providers from noisy reputations. In *Proceedings of the Second Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 1114–1115, New York, NY, 2003. ACM Pres.

[14] H. A. Simon. *Models of Man*. John Wiley & Sons, New York, 1957.

[15] J. Stimpson, M. A. Goodrich, and L. C. Walters. Satisficing and learning cooperation in the prisoner's dilemma. In B. Nebel, editor, *Proceedings of the seventeenth International Conference on Artificial Intelligence (IJCAI-01)*, pages 535–544, San Francisco, CA, Aug. 4–10 2001. Morgan Kaufmann Publishers, Inc.

[16] M. Tennenholtz, Y. Shoham, and A. Schaerf. Adaptive load balancing: A study in multi-agent learning. *Journal of Artificial Intelligence Research*, 2:475–500, Nov 1995.

[17] R. Tietz and O. Bartos. Balancing of aspiration levels as fairness principle in negotiations. Lecture note in *Economics and Mathematical Systems 213*, 1983.

[18] H.-J. Weber. Theory of adaptation of aspiration levels in a bilateral decision setting. In *Zeitschrift für die gesamte Staatswissenschaft*, pages 582–91, 1976.

[19] B. Yu and M. P. Singh. Searching social networks. In *Proceedings of the Second Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 65–72, New York, NY, 2003. ACM Pres.

[20] B. Yu, M. Venkatraman, and M. Singh. An adaptive social network for information access: Theoretical and experimental results. *Applied Artificial Intelligence*, 17(1):21–38, January 2003.