Learning to Use Referrals to Select Satisficing Service Providers

Teddy Candale and Sandip Sen

Mathematical & Computer Sciences Department
University of Tulsa
{teddy-candale,sandip}@utulsa.edu

Abstract. We investigate a formal framework where agents use referrals from other agents to locate high-quality service providers. Agents have common knowledge about providers which are able to provide these services. The performance of providers is measured by the satisfaction obtained by agents from using their services. Provider performance varies with their current load. We assume that agents are truthful in reporting interaction experiences with providers and refer the highest quality provider known for a given task. The referral mechanism is based of the exchange value theory. Agents exchange both the name of the provider to use and the satisfaction obtained by using a referred provider. 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. We experimentally demonstrate the effectiveness of our algorithm in producing stable system configurations where reasonable satisfaction expectations of all agents are met.

1 Introduction

Agents searching for high-quality service providers can either rely on their own interaction experience or use referral services, or referrals from other agents. We assume that the performance of a provider depends both on its intrinsic characteristics and the current workload it is handling. There can also be some intrinsic variabilities in the performance of a provider for the same given workload. A key question to consider in referral systems is the motivation for referring a resource or a provider to another agent. Agents need to find the best tradeoff between helping others with good referrals where the expectation is to get quality referrals in exchange and its potential long-term loss due to increased usage of the referred provider by the requesting agent.

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

M.G. Hinchey et al. (Eds.): WRAC 2005, LNAI 3825, pp. 13–22, 2006.

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, both choosing service providers and making referrals, on its environment.

As such it might not be feasible to seek or determine strategies for optimizing performance. Rather, we posit that agents are concerned about finding a quality of service which exceeds an acceptable performance threshold, γ. This formulation is consistent with Simon and others view of bounded rationality of decision makers within the context of complex organizations [2,7]. Other approaches from game theory also use the notion of *aspiration levels* to stabilize systems and reach equilibrium [1,8,9,10,12].

In this paper, we provide an approach for trading referrals using which agents can locate high-quality service providers. Our proposal involves learning to rate referrers and use such ratings to adjust future referrals to identify effective service providers. We experimentally demonstrate the convergence to satisfactory service provider selections for the entire group of peer agents.

2 Framework

We now present a framework for the environment and the algorithm representing the decision procedure used by the agents to seek and use referrals about service providers.

2.1 Environment

Each agent is assigned a total load of L for each of several different task types per day. At the outset, each agent knows the set of providers that can process each task type without the knowledge of their intrinsic capabilities, or their performance functions. The agents are also unaware of the current service load of any provider.

We designate $\mathcal{E} =< \mathcal{A}, \mathcal{T}, \mathcal{P} >$ as the environment, where

- $\mathcal{A} = \{a_k\}_{k=1..K}$ is the set of agents,
- $\mathcal{T} = \{t_n\}_{n=1..N}$ is the set of task types,
- $\mathcal{P} = \bigcup_{n=1}^{N} \mathcal{P}_n$ where $\mathcal{P}_n = \{p_n^m\}_{m=1..M_n}$. M_n is the number of service providers for task type n. p_n^m is a provider which can perform a task of type t_n.

A provider, p_n^m, is characterized by a *performance* function, f_n^m, which models its performance, e.g., the task processing quality. This function should satisfy the following conditions:

- $f_n^m(0) = 0$
- $f_n^m(+\infty) = +\infty$[1]
- $f_n^m \nearrow$[2]

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 satisfaction levels for the same quality of performance. We represent the satisfaction of an agent by a *subjective* function S_{k,t_n} which models the satisfaction obtained by interacting with any provider of a particular task type. To be more precise, if *perf* is the performance of provider after having performed a task of type t_n then the satisfaction of the agent a_k is $S_{k,t_n}(perf)$.

2.2 Interaction with Other Agents

When an agent needs to process a task and is not cognizant of a satisfactory service provider for the corresponding task type, it approaches other agents for referrals. The model of interaction between agents is inspired by Piaget's theory of exchange value [3]. The asking agent a_a gives its load L and the type of the task t it has to process, which are the *real value*s exchanged. The helping agent a_h answers by providing a name of a provider p and its estimation e of its quality (*real value*). Depending on the received referral with quality estimates, and its evaluation of the reliability or trustworthiness of the helping agent, the asking agent may or may not use the referral. If it decides to use that referral, it asks the referred provider to perform its task t with the load L. At the end of that day, the provider responds with a value, *perf*, representing the quality of service. *perf* depends on the provider's load for the day and on its performance function, f. More precisely, if L_t is the sum of all loads on to this provider then $perf = f(L_t)$. The asking agent computes the satisfaction s it gets given *perf* by using its satisfaction function S_{k,t_n} ($s = S_{k,t_n}(perf)$). Finally, this agent communicates this *virtual value* to the helping agent. a_a regards s as a debit to a_h and a_h regards s as credit from a_a. An agent expects others to help it in return when it has previously done. Consequently, it may have the incentive to ask for help those ones it has helped, i.e., those for which it has a lot of credits. In this model, we assume agents are truthful in reporting their satisfaction obtained from using a referred provider.

In addition to finding quality service providers via referrals, agents also benefit from learning about the quality of referrals provided by peer agents. Agents keep track of their previous interactions with others in order to identify agents who have helped, i.e., provided high-quality referrals. Correspondingly, an agent may be inclined to help those who have helped in the past and less inclined to provide referrals to others. Besides, the information recorded about referrers can be useful to estimate the quality of a provider.

[1] $f_n^m(+\infty) = +\infty$ denotes $\lim_{x \to +\infty} f_n^m(x) = +\infty$.

[2] $\nearrow$ ($\searrow$) means a function is increasing (decreasing).

2.3 Information Recorded by an Agent

An agent a_k will keep track of the following information from its interaction with others where S_i is the satisfaction obtained on day d_i:

- $\Gamma_n^m = [\,(S_i, d_i)\,]$ list of satisfactions obtained by the agent a_k from providers p_n^m,
- $C_{kk'} = \langle [\,(S_i, d_i)\,] \rangle$ is a vector of lists of its credit with agent $a_{k'}$.
- $D_{kk'} = \langle [\,(S_i, d_i)\,] \rangle$ is a vector of lists of its debit with agent $a_{k'}$.
- $\Delta_{kk'} = \langle [\,(S_i, d_i)\,] \rangle$ is a vector of lists of the difference between the actual satisfaction obtained from provider referred by $a_{k'}$ and the estimation agent $a_{k'}$ gives for that provider. This information will be useful for a_k to evaluate future referrals.

The $C_{kk'}$, $D_{kk'}$, and $\Delta_{kk'}$ vectors are indexed by task types, i.e., the jth element of such a vector contains data about interactions involving tasks of type j.

2.4 Satisfaction Criteria

As we have argued previously, in a distributed, open, dynamic environment, it might be more meaningful to seek satisficing rather than optimal performance. This is particularly true in our framework as the number of service providers, their load and hence quality, the peer referring agents as well as their estimations of different service providers can all vary over time. We present a satisficing approach to service provider selection where agents will not consider switching from a service provider if the latter's performance is above a pre-specified threshold or aspiration level. We say that an agent a_k is satisfied by the performance of a provider with performance higher than a threshold $\gamma_{k,n}$ for tasks of type t_n.

2.5 Evaluating a Referrer or a Provider

We present a recency-biased weighting scheme on past interactions to estimate the performance of a provider or the referral ability of another agent.

Definition 1 ***(Weighted mean and weighted standard deviation)*** *Let $\mathcal{S} = [\,(S_i, d_i)\,]_{i \in I}$ where S_i is the satisfaction got at date d_i. Let d be the current day. Let ω be a function such that ω is a decreasing function, $\omega(0) = 1$. We call* **mean** *and* **standard deviation weighted** *by ω the values*

$$\begin{cases} \mu_\omega(\mathcal{S}) = \dfrac{\displaystyle\sum_{i=1}^{|I|} \omega(d - d_i) \cdot S_i}{\displaystyle\sum_{i=1}^{|I|} \omega(d - d_i)} \\ \sigma_\omega(\mathcal{S}) = \Big(\displaystyle\sum_{i=1}^{|I|} \big(\omega(d - d_i) \cdot S_i - \mu_\omega(\mathcal{S}) \big)^2 \Big)^{\frac{1}{2}} \end{cases}$$

The function ω determines the emphasis placed by an agent on interactions over time. If $\omega(d) = 1\ \forall d$ then we get the classical statistical measures of unweighted mean and standard deviation. Besides, the quicker ω decreases the closer μ_ω is to recent obtained satisfaction.

Definition 2 (Weighted experience) *Let e be a function such that $e(0) = 1$, $e(+\infty) = 0$ and $e \searrow$. We call* **experience for $\mathcal{S}$**, *denoted $e_{\mathcal{S}} = [\,(S_i, d_i)\,]_{i\in I}$ the value $\sum_{i=1}^{|I|} e(d - di)$.*

Definition 3 (Class of Usefulness functions) *The* **usefulness function** *of an agent a_k (UF), h_k, are functions such that $h_k : [0,\ 1] \times \mathbb{R} \mapsto [0,\ 1]$ (the first argument is the mean of performance, the second is the experience),*
$h_k(m, 0) = 0$, $h_k(m, +\infty) = m$, $h_k(m, \cdot) \nearrow$
$h_k(0, e) = 0$, $h_k(\cdot, e) \nearrow$

The function h_k combines both the rating an agent has about another agent or a service provider and the amount of experience it has with that entity. The more the experience the more reliance the agent places on the ratings. In other words, μ_ω allows an agent to assess others. e is a measure of the amount of confidence an agent may have of its evaluation. Indeed, the more the experience the more reliance the agent places on the ratings.

2.6 Finding an Agent for Seeking Referral

When agent a_k needs a referral it tries to find an agent which has in the past referred high-quality providers and is also willing to provide referrals and have been consistent in returning help-giving behavior. Each agent $a_{k'}$ is assigned a likelihood $\omega_{k'}$ which increases with the quality $q_{k'}$ of previous referral given for the task a_k's task t_n: $q_{k'} = h_k\big(\mu_\omega(D_{kk'}[t_n]), e_{D_{kk'}[t_n]}\big)$. Additionally, a_k will be inclined to help agents which had helped previously. Let $bal(a_k, a_{k'})$ be the difference between all credits and all debits of a_k from $a_{k'}$. $bal(a_k, a_{k'}) > 0$ means a_k has given more information to $a_{k'}$ than it has received from it. An agent will prefer agents with which it has more credit: $\omega_{k'}$ increases with $bal(a_k, a_{k'})$.

2.7 Performing a Task

When agent a_k needs to perform a task, it evaluates the expected satisfaction for a provider p_n^m as $S_n^m = h_k(\mu_\omega(\Gamma_n^m), e_{\Gamma_n^m})$. It will choose the provider with the greatest expected satisfaction if that is greater than γ_k.

If a_k thinks it may not get the satisfaction it expects given the information it has collected previously, it will ask other agents for referral. A referrer $a_{k'}$ is chosen as described above (Section 2.6). $a_{k'}$ is approached by a_k as described in Section 2.2. $a_{k'}$ answers by given an estimation of the quality $q_{k'}$ of the referral it is giving. a_k will correct $q_{k'}$ by using the information it has collected in $\Delta_{kk'}[t_n]$ (defined in Section 2.3). Let ϵ be a number chosen from Gaussian distribution

$\mathcal{N}\big(\mu_\omega(\Delta_{kk'}[t_n]), \sigma_\omega(\Delta_{kk'}[t_n])\big)$. $q_{k'}$ is corrected by adding ϵ to it. a_k corrects the estimation of $a_{k'}$ since the way of evaluating a provider is not necessarily the same between two agents (different satisfaction functions). Besides, it helps also to correct possible deception attempts. The referral is chosen with probability $q_{k'}$. If it is not chosen, another agent will be approached. If no provider has been chosen when all agents have given their referrals, then a referral will be chosen among those previously referred with likelihood $q_{k'}$.

Once a provider is elected, it is asked to perform the task with load L. At the end of the day this provider will return a value of the performance $perf$ with which it performed the task. $perf$ is computed with the performance function f with the sum of all loads ordered during the day as parameter.

2.8 Responding to a Request for Referral

When approached for referral, we assume that the helping agent refers the best service provider known to it for the corresponding task type. Therefore, the agent refers the provider it thinks is the best given its estimation for the expected quality of service. We also limit the effects of chains of referrals on the stability of the system by imposing the restriction that an agent is not allowed to refer a provider it previously received as a referral.

3 Characterizing Satisficing System States

In this section, our objective is to characterize preferable distribution of agents over providers based on agent and provider parameters like S_{k,t_n}, f_n^m, etc.

In this initial study, we restrict ourselves to a specific variant of the general model we have described above:

- the number of providers is constant equal to M over each type
- the load of an agent is the constant, L, for each type and for each day and is equal for all agents
- the satisfaction function S_{k,t_n} and the aspiration level, $\gamma_{k,n}$, is the same for each agent and for each task type and is denoted by γ.

The satisfaction obtained by an agent when using a service provider on a given day depends only on the performance function of that provider and the total load on that provider for that day.

Definition 4 (Distribution of agents over providers) *We call* **distribution of agents over providers** *for a type t_n the set $\mathcal{D}_n = \{\mathcal{A}_m\}_{m=1..M}$ such that:*

- $\mathcal{A}_m \subseteq \mathcal{A}$ *($\mathcal{A}_m$ may be empty)*
- $\displaystyle\bigcup_{m=1}^{M} \mathcal{A}_m = \mathcal{A}$
- $m_1 \neq m_2 \Longrightarrow \mathcal{A}_{m_1} \bigcap \mathcal{A}_{m_2} = \emptyset$

Definition 5 (Set of satisfaction of a distribution) *The set of satisfaction of a distribution of agents over providers (denoted by $\mathcal{S}(\mathcal{D}_n)$) is the set of satisfactions obtained by the agents in that distribution. More formally, $\mathcal{S}(\mathcal{D}_n) \stackrel{def}{=} \{S(|\mathcal{A}_m| \cdot L)\}_{m=1..M}$. Note that this is a set with no duplicate values.*

Definition 6 (γ-acceptable distribution) *A distribution $\mathcal{D}_n$ is called γ-bacceptable if and only if*

$$\forall s \in \mathcal{S}(\mathcal{D}_n) \quad \gamma \leq s$$

A γ-acceptable distribution is a distribution where the satisfaction of each agent from the provider it is currently using is more than its aspiration level. If agents arrive at such a distribution, they will not be willing to change their choice of service providers, and hence the system will be in equilibrium.

4 Experimental Results

We have implemented this framework by choosing the following parameters: 5 different task types, 10 agents, $L = 1$, the performance functions of providers are equal to $(\alpha \cdot x + \beta)^3 - \beta^3$ (we have used $\beta = 1$), and the satisfaction function of agents are equal to $\dfrac{1}{1 + 0.7 \cdot x^2}$.

We have run several sets of experiments by varying γ, the number of providers and the parameter α. We observe if the daily distribution has stabilized after some days and if the final distribution is γ-acceptable. The table 1 shows the set of satisfaction given the following configuration of service providers:

- 3 providers with $\alpha = 0.1$
- 2 providers with $\alpha = 0.12$
- 3 providers with $\alpha = 0.15$
- 4 providers with $\alpha = 1$

We observe that for $\gamma \in \{0.6, 0.7\}$, there is convergence to a γ-acceptable distribution. In the case $\gamma = 0.6$ there exists two γ-acceptable distributions: both of them occur in our runs and there is no particular way to predict which one will be selected. When $\gamma = 0.7$ only one γ-acceptable distribution exist and hence it is the distribution which is chosen by agents. For $\gamma = 0.75$ no stable behavior is observed. In this case, there is no γ-acceptable distribution. So, the system cannot stabilize since when some agents are unsatisfied, they prevent others from being satisfied by continuing to move from one provider to another.

4.1 Influence of γ on the Speed of Convergence

We further evaluate the influence of γ on the speed of convergence to satisfactory distributions. For each value of γ we run 10 experiments and calculate the average number of days necessary to reach the convergence (see Figure 1). We can see that for small values of γ the convergence is reached very quickly. In these cases,

Table 1. $\mathcal{S}(\mathcal{D}_n)$

γ	$\mathcal{S}(\mathcal{D}_n)$ for Type i
0.6	– $i = 1$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$ – $i = 2$ $\{0.93,\ 0.84,\ 0.73,\ 0.63\}$ – $i = 3$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$ – $i = 4$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$ – $i = 5$ $\{0.9,\ 0.84,\ 0.73,\ 0.63\}$
0.7	– $i = 1$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$ – $i = 2$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$ – $i = 3$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$ – $i = 4$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$ – $i = 5$ $\{0.93,\ 0.9,\ 0.84,\ 0.73\}$
0.75	No convergence

the number of γ-acceptable distributions is high and so is the probability of finding one which leads to quick convergence. However, when γ gets close to 0.75 the number day required to find out a distribution where every agent is satisfied increases significantly. This is explained by the fact the number of γ-acceptable distributions become very small in this range: 2 for $\gamma = 0.65$ and 1 for $\gamma = 0.7$.

4.2 A Case with a Continuous Load

We also ran experiments with the same configurations as before but with a daily load which follows a normal distribution with a mean equal to 1 and a standard deviation equal to 0.5. We wanted to compare convergence results in this case with the constant load situation. We ran a 300 day simulation and calculated the average daily load of each provider for the last fifty simulation days. We present in Table 2 the average satisfactions obtained by each agent. We can see that when $\gamma \in \{0.5,\ 0.6,\ 0.65\}$ the average satisfaction is "γ-acceptable". Consequently, we can expect some kind of convergence for the continuous case.

5 Related Work

Referral systems have recently received increasing attention among multiagent researchers. In [11], Yu and Singh study a referral system when an agent helps the human user find relevant expertise and protect him/her from too many irrelevant requests. In our previous work we have studied the use of referrals to locate service providers when an agent first enters a new community with

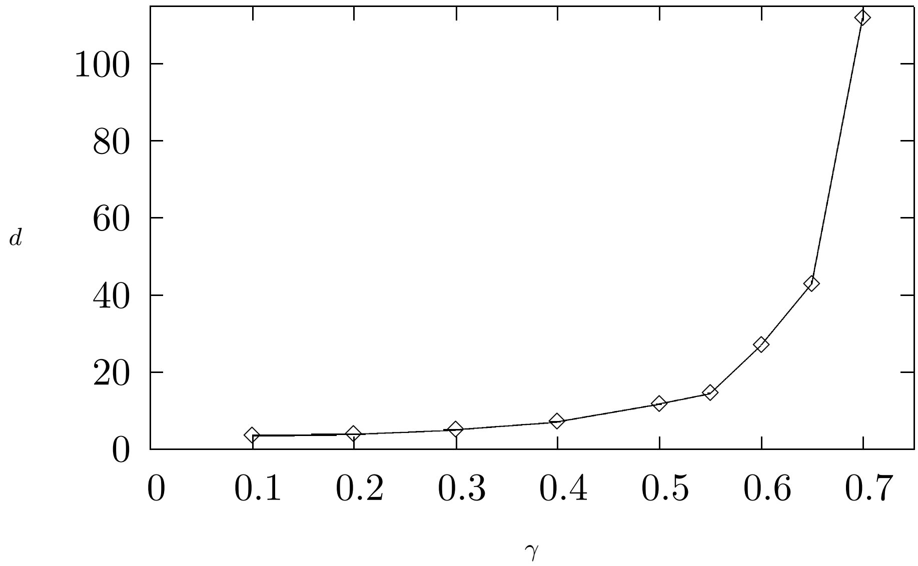

Fig. 1. Effect of γ on the number of days required to reach convergence

Table 2. Satisfaction obtained by agents with a continuous daily load

γ	Average satisfaction
0.5	{0.88 0.77 0.72 1 1 0.90 0.89 0.83 0.83 0.81}
0.6	{0.84 0.80 0.79 0.91 0.97 0.82 0.82 0.83 0.87 0.88}
0.65	{0.85 0.85 0.77 0.81 0.76 0.99 0.99 0.75 0.79 0.87}
0.7	{0.87 0.83 0.83 0.28 0.65 0.83 0.88 0.90 0.83 0.88}
0.8	{0.880.87 0.88 0.065 0.072 0.87 0.89 0.87 0.86 0.82}

no prior knowledge of the quality of service providers or the reliability of the referrers [5,6]. The use of exchange values, motivated by Piaget's theory, for social reasoning in artificial societies is discussed in [4].

6 Conclusion and Future Work

We have presented a framework for agents to choose high-quality service providers by utilizing referrals from other peer agents. We assume that the quality of service provided depends on the intrinsic properties of a provider and the total workload it faces on a given day. Our model accommodates different satisfaction levels for different agents for the same provider performance. We are primarily interested in agents making satisficing choices in that they do not change providers as long as their satisfaction exceeds an *a priori* aspiration level.

Agents need to learn both the quality of service providers as well as the capability of other agents as referrers. Experimental results confirm convergence of agent populations to service providers in satisficing distributions for a large range of aspiration levels, both for constant and varying task loads for agents.

We are currently working on theoretically characterizing requirements for convergence to satisfactory distributions. This paper makes some simplistic assumptions of agent truthfulness while reporting satisfaction to referrers, providing best referrals when asked, always responding to request for referrals, etc. We plan to investigate more complex scenarios involving exploitative and deliberately malicious agents in the population.

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

References

1. G. Kersten. Nego-group decision support system. *Information and Management*, 8:237–246, 1985.
2. J. G. March and H. A. Simon. *Organizations*. John Wiley & Sons, 1958.
3. J. Piaget. *Sociological Studies*. Routlege, London, 1995.
4. M. R. Rodrigues, A. C. da Rocha Costa, and R. H. Bordini. A system of exchange values to support social interactions in artificial societies. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 81–88. ACM Press, 2003.
5. S. Sen and N. Sajja. Robustness of reputation-based trust: Boolean case. In *Proceedings of the First Intenational Joint Conference on Autonomous Agents and Multiagent Systems*, pages 288–293, New York, NY, 2002. ACM Pres.
6. 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.
7. H. A. Simon. *Models of Man*. John Wiley & Sons, 1957.
8. J. L. Stimpson, M. A. Goodrich, and L. C. Walters. Satisficing and learning cooperation in the prisoner's dilemma. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*, pages 535–540, 2001.
9. R. Tietz and O. Bartos. Balancing of aspiration levels as fairness principle in negotiations. In R. Tietz, editor, *Aspiration Levels in Bargaining and Economic Decision Making*. Springer-Verlag, Berlin, 1983.
10. H.-J. Weber. Theory of adaptation of aspiration levels in a bilateral decision setting. *Zeitschrift für die gesamte Staatswissenschaft*, pages 582–91, 1976.
11. B. Yu and M. P. Singh. Searching social networks. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 65–72. ACM Press, 2003.
12. S. Zionts, T. Stewart, and V. Lotfi. An aspiration-level interactive model for multiple criteria decision making. *Computers in Operations Research*, 19(7):671–681, 1992.