Effect of deceptive referrals on system stability

Ikpeme Erete
ikpeme-erete@utulsa.edu
University of Tulsa

Teddy Candale
teddy-candale@utulsa.edu
University of Tulsa

Sandip Sen
sandip@utulsa.edu
University of Tulsa

ABSTRACT

We study the problem of agents attempting to find quality service providers in a distributed environment. While referrals from other agents can be used to locate high-quality providers, referrers may be malicious and provide incorrect referrals to reduce traffic to their preferred service providers. Whereas apparently it would seem that such deceptive referrals can disrupt system stability, we observe that homogeneous groups of deceptive referrals converge faster to stable agent distributions over service providers compared to homogeneous groups of truthful referrers. We conjecture that deceptive referrers can unwittingly reduce the entropy, a measure of volatility, of the system as the recipient of a bad referral may not be inclined to move even if it is not satisfied with its current service providers. Additionally, we observe that mixed groups of truthful and deceptive referrers converge faster to stable distributions compared to homogeneous group of truthful referrers. These results highlight the unexpected positive effect of deceptive agents in stabilizing a population.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Coherence and coordination, Multiagent systems, Intelligent agents*

General Terms

Performance, Experimentation

Keywords

Referral system, deceptive referral, satisficing distribution

1. INTRODUCTION

We investigate the usefulness of referrals from other agents in locating service providers with satisfactory performance where provider performance depends on its workload. In particular, we are interested in quick convergence of such systems to stable agent distributions over the providers where each agent is satisfied with the performance of the provider it is currently using. In previous work, we determined that truthful referral helps increase the speed of convergence as opposed to systems without referrals. But agents may also deliberately provide misleading referrals to, among other reasons, prevent "crowding" at their favorite providers. This paper aims at determining the effect such deception would have on the system convergence and stability.

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, e.g. all factors affecting performance are not directly observable. 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. In this paper, we evaluate the effects of deceptive referrals and if it can, counter-intuitively, mitigate the lack of accurate information available to agents and can maintain or even improve convergence rates to stable configurations.

2. EXPERIMENTAL FRAMEWORK

Our environment is formally defined as follows:
$\mathcal{E} =< \mathcal{A}, \mathcal{R}, perf, L, S, \Gamma >$ 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.
- $f: \mathcal{R} \times \mathbb{R}_+ \rightarrow [0, 1]$, intrinsic performance function of a provider given total load.
- $L : \mathcal{A} \rightarrow \mathbb{R}_+$, load function for the agents.
- $S: \mathcal{A} \times [0,1] \rightarrow [0, 1]$, satisfaction function of agents.
- $\Gamma = \{\gamma_1, \ldots, \gamma_K\}$, the set of satisfaction thresholds of agents.

On each day, agent a_k is assigned a load $L(a_k)$. Initially, each agents know the set of service providers that can handle their tasks but not the intrinsic capabilities, $f(r_n, \cdot)$, of any provider r_n or their respective load capacity. The

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

above formulation is consistent with our previous work on referrals [1].

Our primary objective is for every agent in the system to obtain a satisfaction above its satisfaction threshold, thereby ensuring a stable system. With no incentive to move the system no longer oscillates and we attain an equilibrium state in finite time. Such a distribution is formally characterized as follows:

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

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

3. INFLUENCE OF INERTIA ON CONVERGENCE SPEED

In uncontrolled systems, agents receive highly variable satisfactions. A system at a distribution close to Γ-acceptable distribution will have the tendency to evolve to a worse distribution and vice versa. 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. Indeed, the performance of the system oscillates between desirable values and highly undesirable ones. We believe this problematic effect can be mitigated by increasing the inertia in the system, where the inertia of an agent is its inclination to persist with a service provider it is not satisfied with and is globally measurable by an inverse function of the number of agents moving at any given time, K_{move}.

The work by Candale and Sen [1] showed that when a high number of agents switch resources simultaneously the performance of the system cannot stabilize and varies between high and low values. As a result, agents receive variable feedback from the environment leading to inaccurate learning. They also showed that when the number of agents moving simultaneously remains low, the system stabilizes quickly. We experiment with the following set of agent types:

AgentTRnD: While deciding on which service provider to use an agent can use its past experience and referrals from selected agents. A referral consists of a name of a provider and an estimation of its quality. While picking an agent to ask for referral it selects an agent, with a probability proportional to the quality of past referrals from that agent, from a set of agents whose expected quality of referral is greater than its satisfaction threshold. When an agent, a_k, fails to get an acceptable referral from other agents or from its knowledge of providers, it resorts to systematic exploration of providers with an exploration probability, α_k. The consequence of excessive exploration is a reduction in the inertia of the system and can lead to system instability.

AgentTRwD: This agent is identical to the previous one except the fact that agents always deceive. They refer the same provider as *AgentTRnD*, but alter the true estimation, *es*. More precisely, if *es* is greater than their satisfaction threshold then they report a lower estimate, and otherwise report a higher estimate.

4. EXPERIMENTAL RESULTS

The environment is setup with the exact number of service providers needed to satisfy all the agents in it.We ran experiments with environments having a population of 100 agents each and results averaged over 50runs. Each environment comprised of purely homogeneous agents, truthful and deceptive, and a mixed population of both agents in equal ratio. Agents have identical properties such as coefficient of exploration, satisfaction threshold, and the load function.

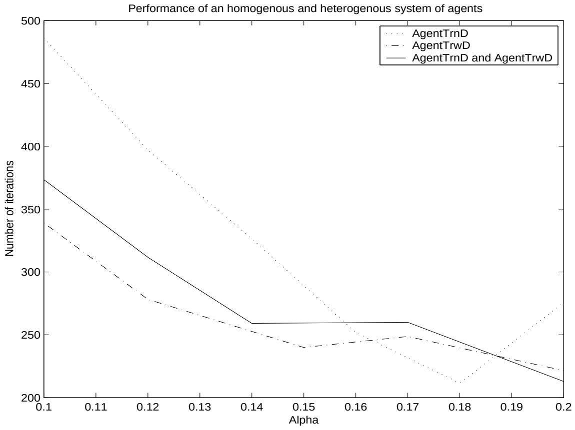

Figure 1: Speed of convergence with varying α: 100 Agents, 50 providers and $\gamma_k = 0.7$.

4.1 Influence of alpha (α)

We observe the speed of convergence of our system by varying coefficient of exploration, α (see Figure 1). We present the following observations:

1. For all agent populations the speed of convergence improves when α increases. For small values of α system convergence is delayed by a slow resource discovery process [1]. For $\alpha > 0.2$, however, exploratory actions become too frequent and increase convergence times. The effect of referrals is more pronounced for smaller values of α. With infrequent exploration, agents have to depend more on referrals to find satisfactory resources.

2. A significant, counter-intuitive result is that the presence of deceptive agents in the population (homogeneous or mixed) leads to better speed of convergence than in homogeneous population of truthful agents. Truthful referrals affect agents using the referred providers, as they will see their satisfactions decrease and then be encouraged to switch resources, which leads to system instability as illustrated in Figure 2. Deception, however, helps an agent maintain the load of its favorite provider at a low level by discouraging probable newcomers. Another key reason is that with deceptive referrals from other agents, it is more likely that an unsatisfied agent may be temporarily led to

believe that it will not receive better satisfaction at other resources. Thus it may choose to not change its provider, thereby decreasing system volatility and ultimately facilitating convergence. Unsatisfied agents, however, will not be persuaded not to change for long, and the system converges steadily.

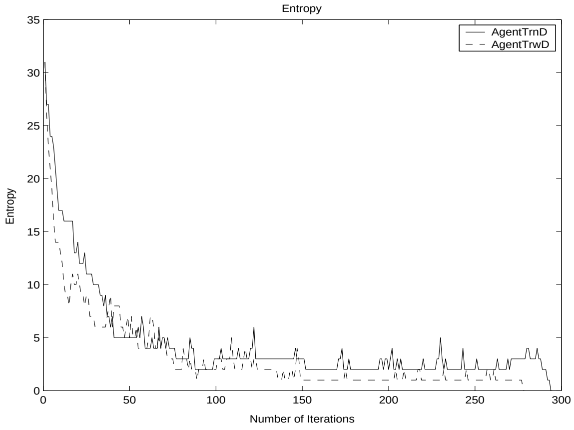

Figure 2: Entropy of truthful and deceptive agents: 100 agents, 50 resources.

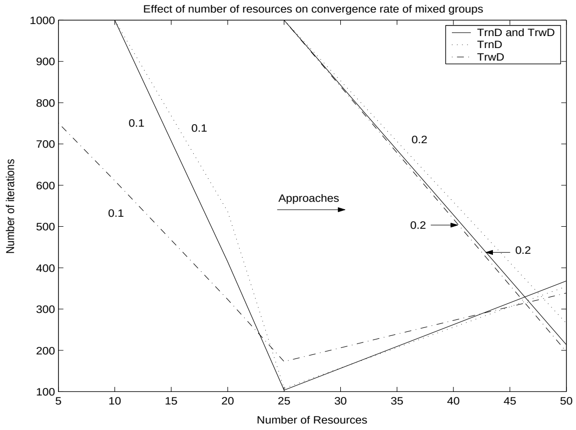

Figure 3: Speed of convergence with varying resources: 100 Agents, $\gamma_k = 0.7,\ \alpha = 0.1,\ ,0.2$.

4.2 Influence of the number of resources

We also investigated the speed of convergence by varying the number of resources (see Figure 3). We observe that deception has a positive effect on the speed of convergence when the numbers of resources is small. With smaller number of resources, it is more likely that agents are dissatisfied with their service providers, leading to greater K_{move}. Deceptive referrals help reduce that for reasons cited above.

5. RELATED WORK

Our domain was inspired by Yu and Singh [4] and Sen,Arora and Roychowdhury [3]. In both work they showed the short term cost of processing referral request and that information can negatively impact agent coordination to find a balanced distribution among resources respectively.

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's [2] study as inertia can be used as a mechanism for controlling K_{move}. Rustogi & Singh [2] also claim that performance is remarkably improved when accepting imprecision.

6. CONCLUSION AND FUTURE WORK

We investigated the effect of the presence of deceptive agents in a population of referrer agents trying to coordinate their selection of service providers. Surprisingly, deception does not harm the system but, in general, helps to enhance the speed of convergence. Agents, by the use of deception, can keep the load of the provider they are currently using at a reasonable level. Besides, deception promotes system stability as it can discourage even unsatisfied agents from moving for some time.

In our future work, we plan on investigating if deceptive agent can still improve system convergence of a truthful referral system with differing satisfaction level and load function. We hope to find ways to achieve an improved convergence in such systems.

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

7. REFERENCES

[1] T. Candale and S. Sen. Effect of referrals on convergence to satisficing distributions. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 347–354,New York, NY, USA,2005. ACM Press.

[2] 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.

[3] 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.

[4] 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.