Using inertia and referrals to facilitate satisficing distributions

Teddy Candale, Ikpeme Erete, and Sandip Sen
{teddy-candale, ikpeme-erete, sandip-sen}@utulsa.edu

University of Tulsa

Abstract. We study the problem of distributed, self-interested agents searching for high-quality service providers where the performance of a service provider depends on its work load. Agents use referrals from peers to locate satisfactory providers. While stable environments may facilitate fast convergence to satisfiying states, greedy and myopic behaviors by distributed agents can lead to poor and variable performances for the entire community. We present mechanisms for resource discovery that involve learning, over interactions, both the performance levels of different service providers as well as the quality of referrals provided by other agents. We study parameters controlling system performance to better comprehend the reasons behind the observed performances of the proposed coordination schemes.

1 Introduction

We study the problem of autonomous agents choosing between several service providers to obtain desired services. We assume a completely distributed environment without central authority or knowledge. Our research goal is to develop mechanisms by which such agent communities can stabilize on states where all agents are satisfied with the service provider they are currently using.

Locating high-quality services is a challenging problem when sharing resources with a large population. Number of service providers are typically limited and their performances depend both on their intrinsic capabilities and workload. Myopic, self-interested behavior can lead to poor performances for the individual and can result in system-wide instability. There is thus a need for non-myopic mechanisms to promote performance and stability of such decentralized systems.

While ideal rational agents may aspire for optimal satisfaction levels, dynamic, partially known, and open environments can render the realization of this ideal improbable. Such an agent is unable to accurately assess the impact of its own decisions, including choice of service providers and making referrals, on the system. As such, it is unrealistic to expect strategies that will always optimize 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 [1, 2, 6, 7].

Referrals from other agents can help agents find more satisfying service providers. But such referrals may cost the referring agent since the load on the referred provider may increase, with corresponding performance deterioration. This is particularly true with referral chains, i.e., if an agent can refer providers it located through referrals from other agents. While referral systems have been widely studied both in theory and in practical applications, 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. 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.

2 Framework

Environment: We present an environment where agents share a set of service providers to perform daily tasks. 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 has a load $L(a_k)$ to perform. a_k assigns its load to a selected provider to handle it in its behalf. At the outset, a_k 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. a_k is also unaware of the current load on the providers. 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 $\mathit{perf} = f(r_n, \sum_{a \in \mathcal{A}_n^d} L(a))$. perf is the service quality received at the end of the day d by every agent in $\mathcal{A}_n^d$. $a_k \in \mathcal{A}_n^d$ will evaluate the performance of r_n by the value $s = S(a_k, \mathit{perf})$. a_k will be satisfied if $s \geq \gamma_k$.

Our aim is to design interaction protocols and behaviors that allow all agents to find satisfying providers. The concept of *distribution* represents how agents distribute themselves over the providers. We call $D = \{\mathcal{A}_n\}_{n=1..N}$ a distribution where $\mathcal{A}_n$ is the set of agents which use provider r_n. A Γ-acceptable distribution is a distribution where every agent is satisfied, i.e, each agent receives a satisfaction above its 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. Consequently, it is an equilibrium concept and our goal is to enable agents to reach such distributions.

Inertia: Oscillations in our environment will happen if at a distribution close to a Γ-acceptable distribution the system has the tendency to evolve to a worse distribution and vice versa. We assume that the total load 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. 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.

An agent may decide to switch resources relying on its own information or on a referrer or to explore to discover either unknown resources or to be able to adapt to changes in the environment. Inertia can be controlled by the following methods:

Exploration: Fast convergence requires learning about provider and referral qualities: more informed decisions will expedite system convergence to satisfactory distributions. Consequently, some systematic exploration of providers is necessary. However, such exploration decreases inertia and can impact convergence rate. An environment where agents explore too much will produce system instability where agents will not have accurate estimations of provider performances since loads vary significantly. In this context, referral systems can be useful since agents may substitute their exploration with others' experiences.

Decision Process: When designing our agents, we chose a "move when you think you can do better"-principle. Consequently, agents never move when they are satisfied. If unsatisfied, agents pick with probability α a resource randomly to ensure exploration. With probability $1 - \alpha$ they try to locate a resource. Henceforth, we refer to the processing in this step as the *decision process*. We present five different decision processes: with and without use of referrals and with more or less inclination to move. We first present decision processes without the use of referral.
NRLI (*No Referral Low Inertia*): This decision process consists in picking a resource for which the agent expects to get at least a minimum level of satisfaction, γ_k^-. Let $es_{k,n}$ be the expected satisfaction agent a_k believes it will get by using resource r_n. Let $\mathcal{R}_{k,\gamma_k^-} = \{r_n \mid \gamma_k^- \leq es_{k,n}\}$ be the set of resources expected to provide satisfaction more than γ_k^-. A resource r_{n_k} is chosen in $\mathcal{R}_{k,\gamma_k^-}$ with likelihood $es_{k,n}$. In the case $\mathcal{R}_{k,\gamma_k^-} = \emptyset$, a_k does not move.
NRHI (*No Referral High Inertia*): NRHI is a variant of NRLI. Agent a_k using NRHI will not move to a provider expected to provide lesser satisfaction than the provider, $r_{n_k^c}$, it is currently using. a_k does not move if $es_{k,n_k} < es_{k,n_k^c}$.
RTLI (*Referral Truthful Low Inertia*): RTLI is also a variant of NRLI. If $\mathcal{R}_{k,\gamma_k^-} = \emptyset$, a_k asks another agent a_{k_h} for referral. a_{k_h} provides both the name of a resource $r_{n_{k_h}}$ and an estimation of the satisfaction it will get ($es_{k_h,n_{k_h}}$). a_k is trustful in the sense it does not try to correct the value $es_{k_h,n_{k_h}}$. a_k will use the referral if $\gamma_k^- < es_{k_h,n_{k_h}}$. Besides, when approached for help, an agent using RTLI is

truthful in the sense that it reports its actual estimate[1]. It refers a resource it would have chosen itself. In other words, it provides the outcome of NRLI.
RTHI (*Referral Truthful High Inertia*): RTHI is a mixture of NRHI and RTLI. When looking for a resource using its own information, an agent uses NRHI and when looking for a referral the agent uses RTLI. When answering a request, it provides the outcome of NRHI.
BRLI (*Balance Referrer Low Inertia*): BRLI is a variant of BRLI. An agent using RFLI will answer a request only from agents with which it has a negative or null balance of exchange. A balance of exchange is the difference between the sum of what it has given and what it has received. More formally, let $bal_{k,k'}$ be the balance maintained by a_k with agent $a_{k'}$. a_k increases $bal_{k,k'}$ by es_{k,n_k} when it provides r_{n_k} as a referral to $a_{k'}$. a_k decreases $bal_{k,k'}$ by $s_{k,n_{k'}}$ where $s_{k,n_{k'}}$ is the satisfaction obtained by a_k if it uses $r_{n_{k'}}$, 0 otherwise.

3 Experimental results

In the previous section, we propose two methods to control the inertia: the coefficient of exploration and the use of decision processes. We will evaluate the two controlling methods while also providing comparisons between referral methods and those without referral.

Experiments comprise a large number, $K = 200$, of identical agents. We use sufficient resources to exactly satisfy the agents present in the environment. In other words, if C_n is the capacity of resource r_n then $\sum_{n=1}^{N} C_n = K \cdot L$ where L is the load imposed by each agent. Hence, we are always sure of the existence of a Γ-acceptable distribution.

We ran experiments to see the influence of the coefficient of exploration α on the speed of convergence. In other words, we measured the number of iterations needed to reach a Γ-acceptable distributionwhen agents use protocols defined in Section 2 given the value of α. Figure 1 presents the result. One environment comprises a high number of resources ($N = 100$) and one comprises a lower number of resources ($N = 40$). We highlight the following observations:

HI performances are much better than those of LI for most values of α. This shows that the speed of convergence is improved greatly if agents decide to move less often. By not moving when it thinks no other resource can satisfy it better than its current resource, an agent avoids conflict of interest since many agents are likely to choose the same resource. Besides, an agent can benefit from the departure of others by staying in its current resource. However, when $N = 100$ and $\alpha \leq 0.02$ both HI and LI have poor results but for different reasons. Detailed analysis of the system given the inertia show us HI have poor performance due to too high inertia; the performance of the system improves very slowly, while LI have poor performances due too low inertia; the system oscillates. In spite of the fact that HI$\leq$LI[2] when $N = 100$ and $\alpha \leq 0.02$, HI is preferrable to LI

[1] In other work, we consider the motivations and the effect for untruthful referrals.
[2] HI$\leq$LI denotes that LI converges faster than HI.

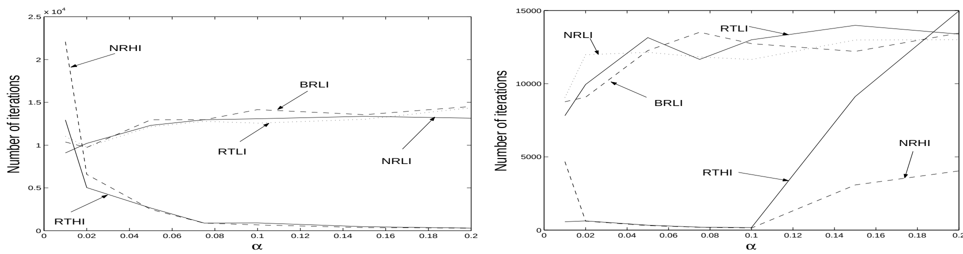

Fig. 1. Number of iterations to reach convergence given α for $N = 100$ (left) and $N = 40$ (right) (200 agents).

because they work better in more environments. Improving HI performance can be achieved more easily by tuning the parameter α.

For high inertia referral decision processes works better, i.e., NRHI $\leq$ RTHI for $\alpha \leq 0.1$. For $\alpha < 0.1$, the inertia is higher with much less exploration, thus preventing substantial improvement of the entropy. The situation is improved by using a referral system. The use of others' information accelerates the resource discovery process. We observe the opposite phenomenon when $\alpha \geq 0.1$, RTHI$\leq$NRHI. With higher values of α, agents are more inclined to explore the environment and hence move more often. This is amplified by the referral system. The use of other's information makes RTHI agents switch resources when NRHI will not.

Performances of LI schemes are equivalent for $N = 100$ and 40. There exists a range of α values in which HI schemes has desirable performances. Detailed studies showed us that for very small values of α the system evolves very slowly with HI since very few agents moves leading to slow convergence. When the values of α are too high, too many agents move simultaneously leading to instable system, i.e., the system oscillate between good states and undesirable states. HI has better scale-up performances when $N = 100$ compare to when $N = 40$.In fact, when the number of resources decreases, assuming the number of agents fixed, more agents have the inclination to move leading to a diminution in the inertia. The range of α values for which HI have desired performances is smaller with lower number of resources.

4 Related work

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 [5]. In their work, 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.

Coordination is a key issue in multiagent systems. Sen et al. [4] show that information can negatively impact agent coordination over resources. They allow agents to move to providers only in the neighborhood of the one they are currently using, thereby achieving perfect coordination faster. They conclude that too much information available to agents lead to oscillating provider loads. This leads to variable provider performances and low convergence speed. Rustogi & Singh [3] study the influence of inertia for system convergence in the same domain. They proved that high inertia speeds up convergence when knowledge increases but low inertia perform better with little knowledge.

5 Conclusion and future work

We have investigated different decision processes to locate satisfactory service providers. These decision processes give agents differing inertia of switching resources given their current and expected satisfications from different resources and can include referrals from other agents. The main conclusion of our experiments is that decision processes with higher inertia of movement (HI procedures) produce faster convergence and better scale-up than those with lower inertia. Even faster convergence with the HI schemes can be produced by using referrals or by tuning the exploration coefficient α. Desirable performances are more difficult to obtain when using LI decision processes regardless of the use of referral systems.

We are currently exploring the effect of non-identical agents and resources. Planned future work includes use of deceptive referral agents and minimalizing such disruptive behavior.

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

References

1. M. Goodrich and J. Stimpson. Learning to cooperate in a social dilemma: A satisficing approach to bargaining. In *International Conference on Machine Learning*, pages 728–735, 2003.
2. J. G. March and H. A. Simon. *Models of Man.* New York, 1958.
3. S. K. Rustogi and M. P. Singh. Be patient and tolerate imprecision: How autonomous agents can coordinate effectively. In *IJCAI '99: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 512–519, San Francisco, CA, USA, 1999. Morgan Kaufmann Publishers Inc.
4. S. Sen, N. Arora, and S. Roychowdhury. Using limited information to enhance group stability. *Int. J. Hum.-Comput. Stud.*, 48(1):69–82, 1998.
5. 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.
6. H. A. Simon. *Models of Man.* New York, 1957.
7. D. S. SJ Russell. Provably bounded-optimal agents. *Journal of Articial Intelligence Research*, 1995.