

Super-Agent Based Reputation Management with a Practical Reward Mechanism in Decentralized Systems

Yao Wang, Jie Zhang*, and Julita Vassileva

Department of Computer Science, University of Saskatchewan, Canada

*School of Computer Engineering, Nanyang Technological University, Singapore
{yaw181}@cs.usask.ca

Abstract. Reputation management is an effective way for agents to find good partners to interact with. In this paper, we propose a novel super-agent based reputation mechanism in decentralized systems. We describe it in the context of web service selection where super-agents with more capabilities maintain reputation information of services and share the information with other consumer agents in the system. A practical reward mechanism is also introduced to create incentives for super-agents to contribute their resources and provide truthful reputation information. We carry out experiments to validate our approach based on a simulated web service environment. The experimental results confirm that our super-agent based reputation management achieves better effectiveness compared to the system that does not use super-agents. The reward mechanism also provides strong support for our super-agent based approach.

1 Introduction

A centralized architecture bears some problems, such as a single point of failure and the need for a high performance server. Moreover, information stored in the central server may become outdated in a dynamic network environment. On another hand, decentralized systems provide good scalability and robustness, a typical example of which is a decentralized agent-based web service system where consumer agents and service providers are distributed across the network [1]. However, such a decentralized environment also creates opportunities for malicious agents to gain benefits without being identified and punished since there is no central authority to monitor them. A consumer agent may have some protection by making a service level agreement (SLA) with providers, which specifies the quality requirement that a received service should meet and the methods of how quality will be measured. SLA usually comes with a cost and requires a third party as an authority to monitor the service.

An alternative approach is to build reputation of services/service providers, a collective measure about quality of the services or honesty of the providers. In this paper, we propose a novel super-agent based reputation mechanism for decentralized systems. In this mechanism, we make use of the resources of super-peers which are agents with more capabilities (i.e. extra CPU power, larger storage and higher network bandwidth). More specifically, super-agents take additional responsibilities: collecting and storing feedback about services, building reputation of services, and sharing the reputation in-

formation with other agents in the system.¹ Consumer agents carefully select which super-agents' advice to follow by modeling the trustworthiness of these super-agents. The consumers that benefit from super-agents also provide the super-agents with feedback about their interactions with service providers after the interaction results are evaluated. We also design a practical reward mechanism inspired by real world examples where service providers offer rewards for agents that bring consumers to consume their services. Super-agents that are honest and contribute more resources will attract a larger number of consumers to follow their advice about services. These super-agents will then be able to obtain more rewards from the service providers. Results of the experiments carried out in a simulated service selection environment confirm that our super-agent based reputation management achieves better effectiveness compared to the system that does not use super-agents. And, with our reward mechanism in place, super-agents are incentivized to contribute more resources and be honest, in order to gain more rewards.

Our idea of using super-agents for reputation management is inspired by studies in P2P (peer to peer) networks. Studies [2] have shown that in practice there is a great heterogeneity in the capability of peers - between three and five orders of magnitude. Peers with poor capabilities become bottlenecks, which degrades the system. With the awareness of the great heterogeneity, pure P2P networks have evolved to super peer networks, such as Kazaa and Gnutella (v0.6) [3]. Super-peers are peers with more capabilities. Peers with poor capabilities are connected to super-peers. Each super-peer acts as a server for a small group of clients (i.e. peers with poor resources) to store their information, and to send and receive messages for them. Super-peers are connected with each other like peers in a pure peer-to-peer system to route, submit and answer queries for their clients and for themselves. Super peer networks work more efficiently than pure P2P networks in terms of searching resources and passing messages.

2 Super-Agent Based Reputation Management

Super-agents play an important role in our super-agent based reputation management. In the context of web service selection, they are also consumer agents but with more CPU processing power, higher network bandwidth, and more availability. Super-agents are responsible for collecting information, building reputation for services, and providing reputation information about services to consumer agents. Super-agents are independent. They decide for which services they want to build reputation. These could be the services that the super-agents are interested in or the services that they want to consume in the future. Some services may have multiple super-agents building reputation for them to avoid possible loss of information and temporary unavailability of some super-agents. When a super-agent wants to build reputation for a service, it tells the service provider to add it into the service's contact list², which contains all super-agents that the service provider knows are building reputation for the service. Super-agents provide reputation information about services upon requests sent by consumer agents.

¹ Although our approach is generally applicable to decentralized systems, we describe it in the context of web service selection in a distributed environment.

² Super-agents do not have to request to be added into services' contact lists. They can also be directly found by consumer agents. Details of this will be provided in the next section.

After gaining help from super-agents, consumer agents also send their feedback about the services to the super-agents. Their feedback helps super-agents build more accurate reputation of the services that will benefit themselves and other consumers.

When a consumer agent wants to find a service, it issues a search query using keywords. If a service provider receives the query, it checks whether it provides the required service. If yes, it returns a message to the querying agent about the service, including the service name and the description of the service. The service provider may also provide a list of super-agents that maintain reputation for the service. If a super-agent receives the search query of the consumer agent, it checks whether it is building reputation for a service matching the search keywords. If yes, it sends a message containing the service's information (e.g. the service name and the description of the service) to the querying agent and also tells the querying agent that it is building reputation for the service. The super-agent may also be asked for reputation value of the service. Based on the received reputation information, the consumer can then model the trustworthiness of the service.

2.1 Trustworthiness of a Service

When a consumer agent c judges the trustworthiness of a service s , it will first use its own experience. After using the service, the agent evaluates the service based on QoS (Quality of Service), which may involve several different metrics, such as response time, accuracy, and reliability. The W3C group provides a summarized guide about defining QoS and its metrics [4]. The overall evaluation of an interaction between a consumer and a service is a combination of the evaluation for each quality metric related to the interaction. How to combine the evaluations of each quality metric depends on the application and a consumer's requirement. For example, for a service providing weather forecast, some consumers may consider the accuracy of the forecast more important. Others may care more about the response time of the service. The result of the overall evaluation about an interaction with the service is either "satisfying" or "not satisfying", which is used to update the consumer agent's trust in the service after the interaction according to the following reinforcement learning formula [5]:

$$T_c(s) = \alpha T'_c(s) + (1 - \alpha)e(s) \quad (1)$$

where $T_c(s)$ denotes the trust value of the service after the update based on the consumer agent's personal experience; $T'_c(s)$ denotes the trust value before the update. α is the learning rate - a real number in the interval (0,1). $e(s)$ is the evaluation result of the interaction represented by a value of 0 for "not satisfying" and 1 for "satisfying".

If the consumer agent c does not have enough personal experience with the service s , it will ask super-agents for information about the service³. The agent sorts the list of super-agents according to its trust⁴ in them from high to low. If the agent's trust in a super-agent is higher than a threshold, the super-agent will be regarded as trustworthy and will be asked for advice about the service. A super-agent which is asked for advice

³ Whether a consumer agent has enough experience with a service is determined by the value of w calculated using Equation 5. If $w = 1$, the consumer agent has enough experience with the service. If $w < 1$, it does not have enough experience.

⁴ The modeling of the trustworthiness of super-agents will be described in Section 2.2.

will provide a reputation opinion about the service, a value in the interval $[0,1]$ where 0 means that the service is totally disreputable and 1 means that the service is completely reputable. Once the consumer agent receives all reputation opinions about the service from all trustworthy super-agents $\{sp_1, sp_2, \dots, sp_n\}$, the consumer agent will calculate an aggregated reputation value according to the following weighted average formula:

$$R_{sp}(s) = \frac{\sum_i^n T_c(sp_i)R_{sp_i}(s)}{\sum_i^n T_c(sp_i)} \quad (2)$$

where $T_c(sp_i)$ is the consumer agent's trust in the super-agent sp_i , and $R_{sp_i}(s)$ is the reputation opinion about the service s provided by sp_i .

The trustworthiness of service s is determined by both the consumer agent's personal experience and the opinions provided by super-agents. It can be calculated based on the combination of the consumer agent's trust $T_c(s)$ in the service based on its own experience and the aggregated reputation opinion $R_{sp}(s)$ of super-agents, as follows:

$$T(s) = wT_c(s) + (1 - w)R_{sp}(s) \quad (3)$$

where w represents how much weight should be put on $T_c(s)$. It is determined based on the number of interactions between the consumer agent c and the service s . We first determine the minimum number of interactions needed for c to be confident about the trust value it has of s computed based on c 's personal experience. Based on the Chernoff Bound theorem [6], the minimum number of interactions can be determined by an acceptable level of error and a confidence measurement as follows:

$$N_{min} = -\frac{1}{2\varepsilon^2} \ln \frac{1 - \gamma}{2} \quad (4)$$

where ε is the maximal level of error that can be accepted by c , and γ is the confidence measure. If the total number of interactions N_{all} is larger than or equal to N_{min} , consumer c has enough personal experience with the service and will be confident about the trust value estimated based on its personal experience. Otherwise, the consumer agent will then also consider the aggregated reputation $R_{sp}(s)$ of the service calculated based on reputation opinions provided by super-agents. The weight w can be measured as follows:

$$w = \begin{cases} \frac{N_{all}}{N_{min}} & \text{if } N_{all} < N_{min}; \\ 1 & \text{otherwise.} \end{cases} \quad (5)$$

When w equals 1, the trustworthiness of the service is the same as the trust value calculated based on only the consumer's personal experience with the service (see Equation 3). When w is less than 1, the aggregated reputation $R_{sp}(s)$ of the service also plays a role in the calculation of the trustworthiness of the service.

Note that there may be the case where a consumer agent does not have enough experience with a service and it also cannot find trustworthy super-agents to ask for reputation opinions about the service. In this case, the consumer agent will also ask advice about the service from other consumer agents which may have interacted with the service before. The calculation of an aggregated reputation value based on other consumer agents' advice and the equation for combining the consumer agent's own experience with the aggregated reputation value are also similar to Equations 2 and 3.

2.2 Trustworthiness of a Super-Agent

The trustworthiness of super-agents is calculated to determine which super-agents will be asked by the consumer agent c for reputation opinions about the service s . It is also used to determine how much weight should be put on each super-agent's reputation opinion in Equation 2.

When the consumer agent asks a super-agent sp_i for a reputation opinion about the service, it can develop trust in the super-agent based on its experience of using the service. A reputation opinion about the service provided by the super-agent $R_{sp_i}(s)$ is represented as a value between $[0, 1]$ calculated based on the super-agent's collected information about the service. After using the service, the consumer agent can evaluate its experience $e(s)$ as "satisfying" or "not satisfying" (1 or 0 respectively). Another reinforcement learning formula is used to model the trustworthiness of the super-agent, as follows:

$$T_c(sp_i) = \alpha T'_c(sp_i) + (1 - \alpha)e(sp_i) \quad (6)$$

where $T_c(sp_i)$ denotes the consumer agent's trust in the super-agent sp_i after the update, and $T'_c(sp_i)$ denotes the trust value before the update. $e(sp_i)$ is the evaluation of the consumer agent's current experience with the advice provided by the super-agent sp_i about the service. It is a real number in the interval $[0,1]$, where 0 means that the reputation opinion about the service provided by the super-agent is false, while 1 means that the reputation opinion is true. $e(sp_i)$ is determined based on the reputation opinion $R_{sp_i}(s)$ as follows:

$$e(sp_i) = \begin{cases} R_{sp_i}(s) & \text{if } e(s) = 1; \\ 1 - R_{sp_i}(s) & \text{if } e(s) = 0. \end{cases} \quad (7)$$

To explain, the value of $e(sp_i)$ is determined by comparing the consumer agent's own experience of using the service, $e(s)$, with the reputation opinion about the service provided by the super-agent. If the consumer agent's experience of using the service is satisfying ($e(s) = 1$), $e(sp_i)$ is equal to the reputation value provided by the super-agent about the service, which is $R_{sp_i}(s)$. If the consumer agent's experience of using a service is not satisfying ($e(s) = 0$), $e(sp_i)$ equals $1 - R_{sp_i}(s)$. For example, if the reputation value of a service provided by a super-agent is 0.9 and the consumer agent's experience is satisfying, the super-agent's reputation opinion is consistent with the consumer agent's experience with the service. In this case, $e(sp_i)$ equals 0.9. However, if the super-agent's reputation opinion is 0.9 and the consumer agent's experience with the service is not satisfying, it indicates that there is a mismatch between the super-agent's opinion and the consumer agent's own experience. Therefore, $e(sp_i)$ equals 0.1. A super-agent can gain more trust if the reputation value it provides matches more closely the consumer agent's experience. The initial value of a consumer agent's trust in a super-agent may be set to 0.5, which means that the super-agent is neither trustworthy nor untrustworthy.

3 Reward Mechanism

We design a reward mechanism to create incentives for super-agents to contribute resources and share truthful reputation information about services. Inspired by real world

examples, our reward mechanism is designed to be rather simple but practical. More specifically, in the reward mechanism, service providers will provide rewards to super-agents. Each service provider can issue its own “virtual points”. This idea is similar to “store credits” in the real world. When a customer accumulates enough “store credits”, these credits can be used to redeem goods in the store. For each consumer agent that consumes a service provided by a service provider, the consumer agent will also tell the provider a list of trustworthy super-agents that have provided advice about the service. To keep our reward mechanism simple, we assume that the “virtual points” will be equally distributed among the trustworthy super-agents providing advice about the service. Service providers in our system have obvious incentives to provide rewards to super-agents. Super-agents building reputation for services offered by the service providers will help the service providers propagate their service information and therefore potentially bring them more consumers. For super-agents, if they build reputation for good services, they can gain “virtual points” from the providers of these good services. The super-agents can then redeem the points for their future interactions with the service providers, i.e. consuming the good services.

Zhang and Cohen [7] have also proposed a trust-based incentive mechanism for encouraging buying agents to truthfully share information about selling agents with other buying agents in electronic marketplaces. Honest buying agents will become neighbors of many other buyers and be rewarded more by selling agents. Their mechanism relies on a trusted central server to compute the reputation of buying agents based on a social network of buyers and to share this information with selling agents, thus scalability becomes a big concern. Our super-agent based reputation management and the practical reward mechanism, on another hand, fit naturally in decentralized networks and make good use of capabilities of super-agents distributed across the networks.

4 Evaluation

To evaluate our super-agent based reputation management and the reward mechanism, we simulate a service selection environment involving service providers and consumers, some of which are super-agents. Consumer agents and super-agents both consume services provided by service providers. A matrix with 5×5 cells is used to simulate the positions of these agents. Originally, service providers, consumer agents and super-agents are randomly located in the cells. Consumer agents and super-agents are different in their ability in discovering service providers. Consumer agents can only find directly the service providers in their own cell. Super-agents are able to directly find the service providers not only in their own cells, but also in the cells adjacent to their own cells. This simulates that super-agents have more searching power than ordinary consumer agents in the network. Super-agents build reputation for services provided by the service providers within their searching scope. In our simulation, super-agents also connect with the consumer agents in their own cells as well as the cells adjacent to them. In this way, consumer agents are able to find through super-agents the service providers that are not in the consumers’ own cells.

There are three types of service providers in our simulation, superior service providers, normal service providers and bad service providers. They are different in their ability

of providing good services. The superior service provider has 90% of probability to provide good services. The normal service provider has 60% of probability, and the bad one has only 40% of probability. The numbers of the three types of providers are 1, 3, and 1 respectively. For the purpose of simplicity, all the services of providers have the same functionality. Each service provider provides one service.⁵ Our experiments involve 100 consumer agents, 10 of which are super-agents. There are 4000 interactions in each experiment. In each interaction, a consumer agent selects a service provider and uses its service. We run each experiment for 10 times and present the average of the results produced by each experiment.




Fig. 1. (a) Successful Interaction Ratio; (b) Honest Super-Agents vs. Dishonest Ones

The goal of the first experiment is to show the effectiveness of our super-agent based reputation management system. We compare the performance of two systems. One system uses super-agents to manage reputation for services provided by service providers and share the results with other consumer agents. The other system does not involve super-agents. We measure the performance of a system by two metrics. One is the ratio of successful interactions. A successful interaction means that a consumer agent selects a service provider, uses its service, and finds it satisfying. By using this measure, we can find out whether super-agents can actually help consumer agents find good service providers. Figure 1(a) shows the ratio of the number of successful interactions over the total number of interactions. From this figure, we can see that our super-agent based system performs much better than the system without super-agents. Super-agents in our mechanism can help consumer agents find potentially good service providers.

An important purpose of our reward mechanism is to create incentives for super-agents to provide truthful reputation information about services. In this experiment, we involve some super-agents that are dishonest. We measure the total number of virtual credits that a super-agent can gain when it acts honestly and dishonestly respectively. As shown in Figure 1(b), honest super-agents can gain many more virtual credits than dishonest super-agents. Our reward mechanism provides strong incentives for super-

⁵ Although we do not consider more complex scenarios in our current evaluation, we argue that a variety of services in simulations will produce the same results to support our proposed approach.

agents to honestly share their reputation information about services with other consumer agents in the system.

5 Conclusion and Future Work

In this paper, we proposed a mechanism of using super-agents to manage reputation of services. Super-agents can build a general public reputation opinion about services and share the reputation information with other consumer agents. A practical reward mechanism is also designed to encourage super-agents to contribute their resources and truthfully share their reputation information. Super-agents can gain rewards from service providers for their contributions and honest behavior. As other existing trust and reputation mechanisms in decentralized systems do not consider the role of super-agents and cannot take advantage of the extra power of super-agents, our research fills the gap and holds good promise when there are more and more super-agents emerging in the networks, which is becoming the reality with the advance of technology, easy access of internet, and lower price for high-performance computers.

Reputation of services built by super-agents reflects a majority consumer agents' opinion. However, agents are different in their interests and judging criteria. The majority opinion may not fit every agent's needs. It would be more desirable to have opinions from agents that are similar in interests and judging criteria. For future work, we will look into the idea of forming communities [8] by super-agents to bring together similar agents that will provide community members with more valuable information about services.

References

1. Han, W.: Integrating peer-to-peer into web services. Master's thesis, University of Saskatchewan, Canada (2006)
2. Saroiu, S., Gummadi, P.K., Gribble, S.D.: A measurement study of p2p file sharing systems. In: Proceedings of Multimedia Computing and Networking (MMCN). (2002)
3. Yang, B., Garcia-Molina, H.: Designing a super-peer network. In: Proceedings of IEEE International Conference on Data Engineering. (2003)
4. Lee, K., Jeon, J., Lee, W., Jeong, S.H., Park, S.W.: Qos for web services: Requirements and possible approaches, world wide web consortium working group note. <http://www.w3c.or.kr/kr-office/TR/2003/ws-qos/> (November 2003)
5. Wang, Y., Vassileva, J.: Bayesian network trust model in peer-to-peer networks. In: Proceedings of IEEE International Conference on Web Intelligence (WI). (2003) 372–378
6. Zhang, J., Cohen, R.: Trusting advice from other buyers in e-marketplaces: The problem of unfair ratings. In: Proceedings of the Eighth International Conference on Electronic Commerce (ICEC'06). (2006)
7. Zhang, J., Cohen, R.: Design of a mechanism for promoting honesty in e-marketplaces. In: Proceedings of the 22nd Conference on Artificial Intelligence (AAAI). (2007) 1495–1500
8. Wang, Y., Vassileva, J.: Trust-based community formation in p2p file sharing networks. In: Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence. (2004)