A Scalable Probabilistic Approach to Trust Evaluation

Xiaoqing Zheng, Zhaohui Wu, Huajun Chen, and Yuxin Mao

College of Computer Science, Zhejiang University, Hangzhou 310027, China
{zxqingcn, wzh, huajunsir, maoyx}@zju.edu.cn

Abstract. The Semantic Web will only achieve its full potential when users have trust in its operations and in the quality of services and information provided, so trust is inevitably a high-level and crucial issue. Modeling trust properly and exploring techniques for establishing computational trust is at the heart of the Semantic Web to realize its vision. We propose a scalable probabilistic approach to trust evaluation which combines a variety of sources of information and takes four types of costs (operational, opportunity, service charge and consultant fee) and utility into consider during the process of trust evaluation. Our approach gives trust a strict probabilistic interpretation which can assist users with making better decisions in choosing the appropriate service providers according to their preferences. A formal robust analysis has been made to examine the performance of our method.

1 Introduction

Just as the Internet is shifting its focus from information and communication to a knowledge delivery infrastructure, the Semantic Web extends the current Web to enable Web entities (software agents, users and programs) to work in cooperation in which information and services are given well-defined meaning. The philosophy of the Semantic Web is that anybody can produce information or provide services or consume and enjoy anyone else's information and services on open environment full of uncertainty and dynamic. There is likely to be a service-rich environment, necessitating the selection between similar services being offered by different providers. The Semantic Web, conceived as a collection of agents, brings new opportunities and challenges to trust research. One of these challenges is modeling trust properly and exploring techniques for establishing computational trust and determining the provenance and quality of content and services. We need to face this important issue of how to decide how trustworthy each information source is and which service we should choose according to these trustworthiness. Trust is a response to uncertainty and uncertainty can be considered as the lack of adequate information to make a decision. Uncertainty is problem because it may prevent us from making the best decision and may even cause a bad decision to be made. Some fundamental questions should be answered before trying to find the best way of modeling trust. These are: what is the exact meaning of trust from the Semantic Web point of view, what information is relevant when evaluating trust, and how to combine information from various sources to produce final trust value. Scalable

K. Stølen et al. (Eds.): iTrust 2006, LNCS 3986, pp. 423 – 438, 2006.

probabilistic approximation seems a direction for future research to deal with this uncertainty. In this paper, we try to answer these questions and propose a composite trust model based on Bayesian sequential analysis which gives trust a formal probabilistic interpretation. Our model combines prior and reputation information to produce a composite assessment of an agent's likely quality and balances the costs (operational, opportunity service charge and consultant fee) and benefits (utility) when communicating or dealing with other agents.

Consider a scenario in which a user (*initiator* agent) want to find a service provider (*provider* agent) to fulfill a special task on the Semantic Web, and his problem is which provider may be the most suitable for him. Assuming that he maintains a list of acquaintances or neighbors (*consultant* agents), and gives each acquaintance a *reliability* factor that denotes what degree this acquaintance's statements can be believed. Each agent also has such a set of acquaintances. During the process of his evaluating the qualities of different providers and making the decision in selecting the best one among them, he can "gossip" with his acquaintances by exchanging information about their opinions on providers' qualities, termed *statements*. This process can be described as using the strategy of exploiting *transitivity*. The idea of this strategy is that an agent sends a message out to request opinions on the quality of the agent who can provide given service. The network of acquaintances of that agent will then either send back an opinion based on experience, or pass the message onto its acquaintances, many of which will be unknown to the first agent. The aim is to enhance the scope of an agent's knowledge by exploring the network feature of agent communities to bring in information from other, unknown, agents. We call it *reputation* of a given provider agent that integrates a number of opinions from acquaintances and acquaintances of acquaintances. Besides reputation information, we also consider initiator agent's *prior information* that is direct experience from history interactions with the provider agent and the various relationships that may exist between them (e.g. owned by the same organization, relationships derived from relationships between the agents' owners in the real life such as friendship or relatives). And then, the *trust* can be generated by incorporating prior and reputation information in our opinion.

The remainder of this paper is organized as follows. In Section 2, a brief overview of research on trust is presented and Section 3 proposes our closed and open trust models based on Bayesian sequential analysis. Then, in Section 4, we give experimental results that show how our models work across a wide variation of the number of agents, the quality of agent population and the accuracy of the survey in terms of precision and the corresponding costs. The conclusions and future work are summarized in Section 5.

2 Related Work

Given its importance, a number of computational models of trust have been developed in security, e-commerce and multi-agents systems. Probably the most widely used trust models are those on *eBay* and *Amazon* Auctions. Both of these are implemented as a centralized rating system so that their users can rate and learn about each other's reputation. For example, in *eBay*, sellers receive feedback (+1, 0, −1) for their reliability in each auction and their trust is calculated as the sum of those ratings over the last six months. Both approaches are completely centralized and require users to explicitly make and reveal their ratings of others. However, it is questionable if the ratings reflect the trustworthy behavior of sellers, since in the online marketplaces, it is very likely for users to misbehave or trick with malice. The worse is that these systems are not convenient for users to receive a personalized set of trusts according to their preferences.

Social network analysis techniques are used in [J. Golbeck et al., 2003] to measure trust over a Friend of a Friend (FOAF) network, extended with trust relations. This work describes the applicability of social network analysis to the Semantic Web, particularly discussing the multi-dimensional networks that evolve from ontological trust specifications. But this work uses simply function to calculate trust and does not consider more in depth investigation of algorithms for calculating trust. So, it is nearly impossible for application in real world. Moreover, the algorithm for evaluating trust in this method is just heuristic and no reasonable explanation in terms of mathematics.

[S. D. Ramchurn et al., 2003] develops a trust model, based on confidence and reputation, and shows how it can be concretely applied, using fuzzy sets, to guide agents in evaluating past interactions and in establishing new contracts with one another. But this model is rather complex and cannot be easily used in today's electronic communities. The main problem with their approach is that every agent must keep rather complex data structures, which can be laborious and time-consuming. Also, it is not clear how the agents get needed information and how well the model will scale when the number of agents grows.

[E.M. Maximilien and M.P. Singh, 2003] proposes a centralized agent to measure the reputation of Web services by monitoring and collecting client feedback, and making this information available to other agents. Relying on centralized institutions to measure trust takes the burden off the interactive agents when deciding which agents to trust. However, such systems raise the question of how trustworthy is the sources of their trust information in the first place, and why such trust warehouses should be trusted at all. We argue against these centralized units for measuring trust because of their scalability limitations and the implicit trust measurement mechanisms they adopt.

[Dong Huynh, 2004] presents FIRE, a trust and reputation model that integrates interaction trust, role-based trust, witness reputation and certified reputation to generate a comprehensive assessment of an agent's likely performance. But this work assumes that all agents are honest in exchanging information, uses static parametric model that can not dynamically adjust themselves to the change of environment, and has no learning abilities, so, it can not be used in real open environment.

[Yao Wang and Julita Vassileva, 2003] proposes a Bayesian network-based trust model for a file sharing peer-to-peer application which enables an agent consider its trust in a specific aspect of another agent's capability or in a combination of multiple aspects. According to this model, peers make recommendations to each other, by exchanging and comparing their Bayesian networks. After this comparison, the agents update their trust ratings of each other, depending on whether they share similar preferences, on the assumption that an agent with similar preferences is more likely to give suitable recommendations than others. However, the model's mathematical formulation for the calculation of trust can at best be described as intuitive—without justifications and their experiments just use a very simple naïve Bayesian network, which cannot represent complex relationships. Furthermore, this model is applicable in small-size network and does not scales well to any social network size because maintaining and comparing more complex Bayesian network for each agent will be computationally intractable.

3 Trust Model

The main point of a trust model is to provide a way to operate under uncertainty, not taking too many risks, not missing too many opportunities, not deliberating too long before making commitments. Therefore, before presenting our trust model it is necessary to understand the risks and benefits associated with the trust evaluation. There are many different kinds of costs and benefits an agent might incur when communicating or dealing with other agents and the trust model should balance these costs and benefits. We begin by extending the viewpoint of [K. O'Hara et at., 2004] and discussing four types of costs: operational, opportunity, service charge and consultant fee and continue by introducing utility function that is used to reflect the preferences of agent's owner.

Operational Cost. Operational cost is the expenses of computing trust value. In other words, this is the cost of setting up and operating the whole trust plan. Therefore, the more complex the algorithm is, the higher this cost is expected to be.

Opportunity Cost. Opportunity cost is the lost of missing some possibility of making better decision via further investigation. Generally, the more observations, the lower the opportunity costs.

Service Charge. Service providers differ from each other not only in their qualities but also in their charges of services. Service charge is that will be paid to the selected provider agent who provides fee-based services.

Consultant fee. Consultant fee is incurred when an agent asks the opinions of other agents who (may be professional in given domain) charge for their opinions.

Utility Function. To work mathematically with ideas of "preferences", it will be necessary to assign numbers indicating how much something is valued. Such numbers are called *utilities*, and *utility theory* deals with the development of such numbers. *Utility function* can be constructed to state preferences and will be used to estimate possible consequences of the decisions.

3.1 Closed Trust Model

In our model, the quality of provider agent can be considered to be an unknown numerical quantity, and will represent it by θ (possibly a vector) and it is possible to treat θ as a random quantity with a probability distribution. Consider the situation of an agent *A* try to make an estimate of agent *B'*s trust value. *A* holds a prior information (subjective) of *B*, represented by distribution $\pi(\theta)$ (for either the continuous or discrete case), and request *A*'s acquaintance to give opinions on *B*'s quality. After *A* receives the assessments of *B*'s quality from its acquaintances, *A* takes these statements as sample about θ. Outcome of these sample is a random variable and will be denoted $\boldsymbol{X}$ (Often $\boldsymbol{X}$ will be a vector). A particular realization of $\boldsymbol{X}$ will be denoted x and $\boldsymbol{X}$ will be assumed to be either a continuous or a discrete random variable, with density $f(x|\theta)$. Then, we can compute "posterior distribution" of θ given x, denoted $\pi(\theta|x)$. Just as the prior distribution reflects beliefs about θ prior to investigation in B's reputation, so $\pi(\theta|x)$ reflects the update beliefs about θ after (posterior to) observing the sample x. In other words, the posterior distribution combines the prior beliefs about θ with the information about θ contained in the sample, x, to give a composite picture of the final beliefs about θ. We take the posterior distribution of θ, $\pi(\theta|x)$, as the estimate of B's trust. If we want to take another investigation on B's quality for more accuracy, $\pi(\theta|x)$ will be used as prior distribution for the next investigation instead of original $\pi(\theta)$.

When several similar provider agents exist, A need to decide which one should be selected. At that time, the preferences of agent A's owner should be considered properly to make this decision. Therefore, *utility function* should be constructed for agent A's owner, which represented by $U_A(r)$, to express his preferences, where r represents rewards of the consequences of a decision. Supposing that $\pi(\theta|x)$ is the posterior distribution of provider agent B, the expected utility of function $U_A(r)$ over $\pi(\theta|x)$, denoted $E^{\pi(\theta|x)}[U_A(r)]$, is possible gain of consequence of selecting B. If there are several provider agents can be considered, we simply select one that will result in the most expected utility as decision.

By treating an agent as a node, the "knows" relationship as an edge and remember that trust is an asymmetric relation, a directed graph emerges. To facilitate the model description, agents and their environment are to be defined. To clarify the idea of our trust model, we begin with a simple illustration. Consider the scenario that agent A is evaluating trust value of B and C for being business. The set of all consultant agents that A requests for this evaluation as well as A, B, C can be considered to be a unique society of agents $\boldsymbol{N}$. In our example (see Figure 1), $\boldsymbol{N}$ is $\{A, B, C, D, E, F, G, H, I, J, K\}$ and is called a "closed society of agents" with respect to A.

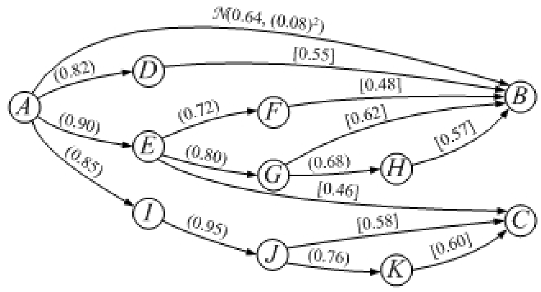

Fig. 1. A "closed society of agents" with respect to agent A

Decisions are more commonly called *actions* and the set of all possible actions under consideration will be denoted $\mathcal{A}$. In our example, initiator agent A is trying to decide whether to select agent B (action b) or C (action c) as business partner

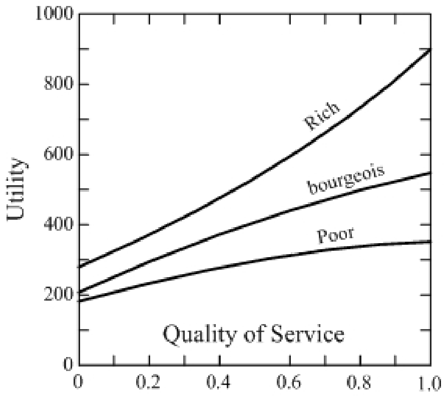

Fig. 2. Utility function

($\mathcal{A}$ = {b, c}). Service charges of B and C are 250 and 210 units respectively (SC_B = 250, SC_C = 210, where SC denotes service charge). We will treat the quality of service, θ, as a continuous variable here, and the unknown quantity θ that affects the decision process is commonly called the *state of nature*. See Figure 1, a notion on an edge between initiator and consultant agent or between the two consultant agents represents reliability factor, that between consultant and provider agent is the assessment of service quality and that between initiator and provider agent is prior information. According to Figure 1, the agent A feels that θ_B, the service quality of B, has a normal prior density, $\mathcal{N}(0.64, (0.08)^2)$ (subscript denotes which an provider agent is being taken consideration). We also suppose that the prior density of θ_C is $\mathcal{N}(0.5, (0.15)^2)$ here[1]. The probability distribution of $\boldsymbol{X}$ that represents the assessments of service quality from consultant agents will, of course, depend upon the unknown state of nature θ. Therefore, we assume that $\boldsymbol{X}$ is another continuous random variable with density $f(x|\theta) \sim \mathcal{N}(\theta, (0.05)^2)$. We also assume that users can be divided into three types, "*Rich*", "*Bourgeois*", and "*Poor*", and agent A's owner belongs to "*Bourgeois*". The curves of their utility functions are shown in Figure 2. We use polynomial regression up to fourth degrees to get fitted model of utility curves and we have

$$U_{Rich}(r) = 280.2 + 334.1r + 705.7r^2 - 944.1r^3 + 524.5r^4$$
$$U_{Bourgeois}(r) = 208.4 + 459.8r - 17.95r^2 - 99.84r^3$$
$$U_{Poor}(r) = 179 + 294.1r - 77.62r^2 - 80.03r^3 + 34.97r^4$$

Note that θ and $\boldsymbol{X}$ have joint (subjective) density

$$h(x,\theta) = f(x \mid \theta)\pi(\theta) \tag{1}$$

and in making decision it is clearly important to consider what the possible states of nature are. The symbol Θ will be used to denote the set of all possible states of nature. So, $\boldsymbol{X}$ has marginal (unconditional) density.

[1] Here, we use simple noninformative prior, [James O. Berger, 1985] discussed other methods to construct prior distribution, such as histogram, relative likelihood, maximum entropy, moment, marginal distribution and ML-Ⅱ approaches.

$$m(x) = \int_{\Theta} f(x \mid \theta)\pi(\theta)d\theta \tag{2}$$

it is clear that (providing $m(x)\neq 0$)

$$\pi(\theta \mid x) = \frac{h(x,\theta)}{m(x)} \tag{3}$$

In discrete situations, the formula for $\pi(\theta|x)$ is commonly known as Bayes's theorem.

Assume a sample $\boldsymbol{X} = (X_1, X_2, \ldots, X_n)$ from a $\mathcal{N}(\theta, \sigma^2)$ distribution is to be taken (σ^2 known), and Let be prior information, $\pi(\theta)$, a $\mathcal{N}(\mu, \tau^2)$ density, where μ and τ^2 are known. Since $\overline{X}$ is sufficient for θ and noting that $\overline{X} \sim \mathcal{N}(\theta, \sigma^2/n)$. Therefore, the posterior distribution θ of give $x = (x_1, x_2, \ldots, x_n)$ is $\mathcal{N}(\mu(x), \rho)$, where

$$\mu(x) = \frac{\sigma^2 / n}{(\tau^2 + \sigma^2 / n)}\mu + \frac{\tau^2}{(\tau^2 + \sigma^2 / n)}\overline{x} \tag{4}$$

$$\rho = \frac{\tau^2\sigma^2}{(n\tau^2 + \sigma^2)} \tag{5}$$

we also need to understand how to calculate the value of n and $\overline{x}$ in the Formula 4 and 5. Following formulas are used to get n and $\overline{x}$, where m represents how much opinions from a variety of sources (sample information) are used to evaluate the trust value. r_i and s_i denote the reliability factor and the statement of service quality for given provider on path i respectively.

$$n = \sum_{i=1}^{m} r_i \tag{6}$$

$$\overline{x} = \sum_{i=1}^{m} \frac{s_i \times r_i}{n} \tag{7}$$

Now, we can evaluate trust value of B for A by using above formulas and information. We take the assessments of B's quality from consultant agents as sample about θ_B and combine these sample information (x) and prior information into posterior distribution of θ_B given x. As shown in Table 1, n_B and $\overline{x}_B$ are:

$$n_B = 0.82 + 0.72 + 0.648 + 0.4896 = 2.6776$$

$$\overline{x}_B = \frac{0.55\times0.82+0.62\times0.72+0.48\times0.648+0.57\times0.4896}{2.6776} = 0.5555$$

Table 1. Calculating n and $\overline{x}$ of Agent B

No.	Path	Statement	Reliability Factor
1	$A\to D\to B$	0.55	0.8200
2	$A\to E\to G\to B$	0.62	0.7200
3	$A\to E\to F\to B$	0.48	0.6480
4	$A\to E\to G\to H\to B$	0.57	0.4896
Total		–	2.6776

Table 2. Calculating n and $\overline{x}$ of Agent C

No.	Path	Statement	Reliability Factor
1	$A\to E\to C$	0.46	0.9000
2	$A\to I\to J\to C$	0.58	0.8075
3	$A\to I\to J\to K\to C$	0.60	0.6137
Total		–	2.3212

hence, $\pi_B(\theta|x) \sim \mathcal{N}(\mu_B(x), \rho_B)$, where

$$\mu_B(x) = \frac{0.05^2/2.6776}{0.08^2+0.05^2/2.6776}\times 0.64 + \frac{0.08^2}{0.08^2+0.05^2/2.6776}\times 0.5555 = 0.5663$$

$$\rho_B = \frac{0.08^2\times 0.05^2}{2.6776\times 0.08^2+0.05^2} = 0.0285^2$$

Note that we employ multiplying to merge two or more than two reliability factors. For example, the reliability factor of the edge $A\rightarrow E$ is 0.9 and that of $E\rightarrow G$ is 0.8, then the value of reliability factor on the path $A\rightarrow G$ is 0.72 (0.9 × 0.8). The reason behind using multiplying is that if the statement is true only if the agents that propagate this statement all tell the truth and it is considered to be independent for any two agents to lie or not to lie. Like B, agent C has a posterior distribution of $\pi_C(\theta|x) \sim \mathcal{N}(0.5370, 0.0321^2)$. After obtaining the final posterior distribution of B and C, we can compare the result of the expectation of $U_A(r)$ over $\pi_B(\theta|x)$ minus SC_B with that of the expectation of $U_A(r)$ over $\pi_C(\theta|x)$ minus SC_C, and simply select the agent that possibly will produce more utility. Expected utility of B and C are (Simpson method is used to solve definite integral):

$$\begin{aligned}\text{Utility of } B &= \int_{-\infty}^{+\infty} U_A(r)\pi_B(\theta \mid x)d\theta - SC_B \\ &= \int_{-\infty}^{+\infty}(208.4 + 459.8\theta - 17.95\theta^2 - 99.84\theta^3)\times \\ &\quad \frac{1}{\sqrt{2\pi}\times 0.0285} e^{-\frac{(\theta-0.5663)^2}{2\times 0.0285^2}} d\theta - 250 \\ &= 194.72\end{aligned}$$

$$\begin{aligned}\text{Utility of } C &= \int_{-\infty}^{+\infty} U_A(r)\pi_C(\theta \mid x)d\theta - SC_C \\ &= \int_{-\infty}^{+\infty}(208.4 + 459.8\theta - 17.95\theta^2 - 99.84\theta^3)\times \\ &\quad \frac{1}{\sqrt{2\pi}\times 0.0321} e^{-\frac{(\theta-0.5370)^2}{2\times 0.0321^2}} d\theta - 210 \\ &= 224.46\end{aligned}$$

hence (224.46 > 194.72), action c should be performed which means that C is more appropriate than B in the eyes of A .

3.2 Open Trust Model

Above discussion is under the condition that the "closed society of agents" must be defined at first, but it is nearly impossible for inherent open and dynamic Web. Our idea is that at every stage of the procedure (i.e., after every given observation) one should compare the (posterior) utility of making an immediate decision with the "expected" (preposterior) utility that will be obtained if more observations are taken. If it is cheaper to stop and make a decision, that is what should be done. To clarify

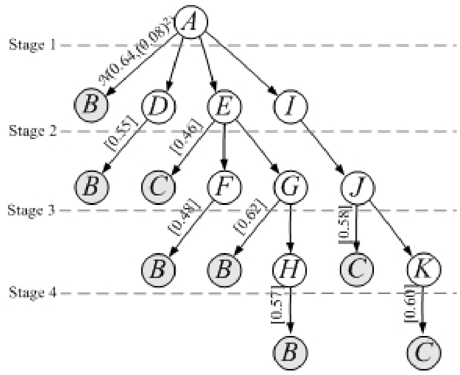

Fig. 3. The process of trust evaluating

this idea, we transform Figure 1 to the structure of tree, shown in Figure 3. The root of tree is an initiator agent, a no leaf node represents a consultant agent (a provider agent also is allowed in real application) and a leaf node represents a provider agent.

The goal of preposterior analysis is to choose the way of investigation which minimizes overall cost. This overall cost consists of the decision loss (opportunity cost) and the cost of conducting observation (consultant fee). Note that these two quantities are in opposition to each other. To lower the decision loss it will generally be necessary to carry out a larger observation, whereby the cost of consultant fee will be increased. In this section, we propose an approach to balance these two costs.

We continue above example used in the illustration of the closed trust model. As shown in Figure 3, we begin at the stage 1 when *A* only hold the prior information of *B* and has no any information about *C* (even the existence of *C*, but it is more likely that an agent with the prior distribution of $\mathcal{N}(0.5, (0.15)^2)$ and the expected service charge of 210 is near in the network). Agent *A* either can make an immediate decision (to select *B*) or can send request to its acquaintances for their opinions by extending the tree of Figure 3 down to the next layer.

Suppose that the cost of consultant service is determined by how much agents will be requested at the next stage and consultant fee is the constant of 1 for each times (for example, at the stage 1, agent *A* can ask Agent *D*, *E* and *I* for their opinions, therefore the cost of consultant fee at stage 1 will be 3). When preposterior and Bayesian sequential analysis are performed in our sample, we use the 1-step look ahead procedure for simplicity.

The utility of an immediate decision is the larger of

$$\int_{-\infty}^{+\infty} U_A(r)\pi_B(\theta)d\theta - SC_B = 217.78$$

here, $\pi_B(\theta) \sim \mathcal{N}(0.64, (0.08)^2)$.

And $\int_{-\infty}^{+\infty} U_A(r)\pi_C(\theta)d\theta - SC_C = 207.54$

here, $\pi_C(\theta) \sim \mathcal{N}(0.5, (0.15)^2)$.

(note that C is not known at this stage, we use subscript of C in above equation just for convenience). Hence the utility of an immediate decision is 217.78. If the request message is sent and x observed, the posterior density $\pi_B(\theta|x)$, is $\mathcal{N}(\mu_B(x), \rho_B)$, where

$$\mu_B(x) = \frac{0.05^2}{0.08^2 + 0.05^2} \times (0.64) + \frac{0.08^2}{0.08^2 + 0.05^2}(x) \cong 0.1798 + (0.7191)x$$

$$\rho_B = \frac{0.08^2 \times 0.05^2}{0.08^2 + 0.05^2} \cong 0.0018$$

However, that we do not know which x will occur, but we know the marginal distribution of $\boldsymbol{X}$, $m(x)$, is $\mathcal{N}(\mu, \sigma^2 + \tau^2)$ and the "predictive" distribution, $m_B(x)$, which in this situation is $\mathcal{N}(0.64, (0.08)^2 + (0.05)^2)$. Note that if $x < 0.5914$ is observed, the expected utility of $\pi_B(\theta|x)$ is less than 207.54, so we prefer to select C instead of B. Hence expected utility of not making immediate decision is

$$\begin{aligned}
&\int_{-\infty}^{0.5914} 207.54\, m_B(x)dx + \\
&\int_{0.5914}^{+\infty} \left(\int_{-\infty}^{+\infty} U_A(r)\pi_B(\theta \mid x)d\theta - SC_B\right)m_B(x)dx - 3 \\
&= 219.24
\end{aligned}$$

This is no other than the opportunity cost (3 is consultant fee in above equation). Because 219.24 > 217.78, and then further investigation would be well worth the money, in other words, A should send request to its acquaintances for their opinions. In order to answer the question of which sample information should be used with higher priority, we prescribe that the sample from the agent with shorter referral distance should be used first.

Remember that the further exploiting should be terminated immediately along the path on which a cycle is detected. Table 3 and 4 show the each stage of the process for agent B and C respectively and the residual process of Bayesian sequential analysis are shown in Table 5.

Table 3. The process of evaluating B's trust

No.	Sources of Opinions	$\mu_B(x)$	ρ_B	Utility
1	None	0.6400	0.0800	217.78
2	*D*	0.5790	0.0454	198.80
3	*D, G, F*	0.5656	0.0311	194.45
4	*D, G, F, H*	0.5663	0.0285	194.72

Table 4. The process of evaluating C's trust

No.	Sources of Opinions	$\mu_C(x)$	ρ_C	Utility
1	None	0.5000	0.1500	207.54
2	*E*	0.4644	0.0497	197.65
3	*E, J*	0.5157	0.0371	216.79
4	*E, J, K*	0.5370	0.0321	224.46

See Table 5, at the stage 3, the expected utility of C begins to larger than that of B, and because 216.79 > 214.78, making an immediate decision is more profitable (There is no need for the opinions of agent H and K). Therefore, A should stop investigating and select C as a decision. The advantage of sequential analysis should be clear now. It allows one to gather exactly the correct amount of data needed for a decision of the desired accuracy.

Table 5. The process of Bayesian sequential analysis

Stage	Agent *B*				Agent *C*				Consultant Fee	Utility		Decision
	Prior Distribution		Marginal Distribution		Prior Distribution		Marginal Distribution			Immediate Decision	Further Investigation	
	μ_B	τ_B	$\mu_{x\mid B}$	$\sigma_{x\mid B}$	μ_C	τ_C	$\mu_{x\mid C}$	$\sigma_{x\mid C}$				
1	0.6400	0.0800	0.6400	0.0943	0.5000	0.1500	–	–	3	217.78	219.24	Continue
2	0.5790	0.0454	0.5790	0.0675	0.4644	0.0497	–	–	3	198.80	199.36	Continue
3	0.5656	0.0311	–	–	0.5157	0.0371	0.5157	0.0623	2	<u>216.79</u>	<u>214.78</u>	<u>Stop</u>

4 Experiments

In this section, we developed a simulation system to measure some properties of our trust models. We present three sets of experiments. The goal of the first experiment is to see if our trust models help users to select the appropriate providers that match better their preferences. We compared the performance of the closed and open models in terms of precision and consultant fee. The second experiment is to examine the effect that varying the accuracy of the survey has on the overall performance of the system, and finally, we want to see what quality of agent population is necessary for the system to work well.

For the sake of simplicity, each agent in our system played only one role at a time, either the role of service provider or the role of consumer (including initiator and consultant agents). The numbers of provider and consumer agent were equal and had half each. Every consumer agent kept two lists. One was the list of "neighbor" that recorded its all acquaintances to each of which a reliability factor was attached. The other was the provider list that recorded the known providers and the corresponding prior information. The number of total items in above two lists is defined as "the degree of outgoing" and in our experiments the degree of outgoing is set to 5.

Following [M. Richardson et al., 2003], we expected the information on the Semantic Web to be of varying quality, so we assigned to each consumer agent i a quality $\gamma_i \in [0, 1]$. A consumer's quality determined what degree that a statement passed or issued by the consumer was true. Unless otherwise specified, the quality of consumers was chosen from a Gaussian distribution with $\mu = 0.8$ and $\sigma = 0.15$. For any pair of consumer i and j where i trust j:

$$t_{ij} = \text{uniformly chosen from } [max(\gamma_j - \delta, 0), min(\gamma_j + \delta, 1)]$$

Where γ_j is the quality of consumer j and δ is a noise parameter that determines how accurate consumers were at estimating the qualities of other consumers, and for these experiments we let $\delta = 0.2$. We also generated randomly each provider agent's quality that was represented by distribution $\pi(\theta)$ (the mean of θ was chosen from a Gaussian distribution $\mathcal{N}(0.5, (0.20)^2)$ and the variant of θ was chosen from a Gaussian distribution $\mathcal{N}(0.08, (0.02)^2)$, see Section 3 for more detail) and its corresponding service charge, and assumed that consultant fee was the constant of 1 for each times. Unless otherwise specified, the probability distribution of $\boldsymbol{X}$ that represents the assessments of service quality from consultant agents (the accuracy of the survey) was $\mathcal{N}(\theta, (0.05)^2)$.

As above mentioned, we assumed that the consumers are divided into three types, "*Rich*", "*Bourgeois*" and "*Poor*", and which type a consumer belong to was decided randomly during the experiments. The curves of the consumers' utility functions are shown in Figure 2.

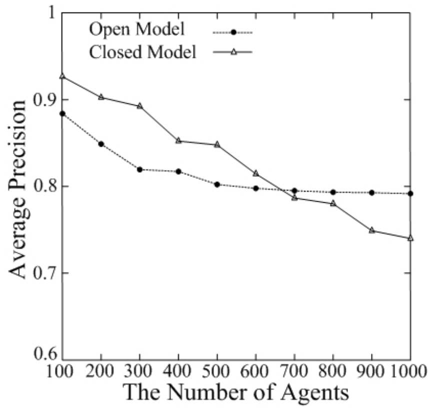

Fig. 4. Effect of the number of agents on the precision

Let G be the set of all provider agents in an experiment. The symbol M_i is used to denote the maximum utility that a provider in G can bring consumer i and let O_i be the utility that is produced by a provider selected by i using certain trust model, so $precision_i$ can be defined as O_i / M_i. The maximum path length was 10 in our experiments. The program would terminate and generate the results when reaching this maximum. We run each configuration for 10 times and use the means for the final experimental results.

Varying the Number of Agents. We explored the effect of varying the number of agents for the closed and open trust models introduced earlier. As shown in Figure 4, we found that the precision differed only slightly between the closed and open models. We also did the experiment for 10000 agents and the average precisions of closed and open models are 0.6147 and 0.7745 respectively. To our surprise, the open model began to outperform the closed model when the number of agents reached to 700. Through careful investigation, we believe this is because the closed model will meet with more noise in the network than the open model when the number of agents grows. We also found that the average precision of the open model decreased slightly when the number of agents grew from 100 to 10000. Therefore, the results show that the open trust model is robust to the population of agents.

As shown in Table 6 and Figure 4, we found that the average consultant fee of the open trust model was significantly lower than that of the closed trust model, though two models differed only slightly in terms of precision. This also meant that the open model had less runtime for trust evaluation. Furthermore, the average consultant fee of the closed model dramatically went up after 200 agents, otherwise, that of the open model increased comparatively smoothly (since the maximum path length was set to 10, the average consultant fee of the closed model did not increase significantly after

300 agents. If the experiments were not restricted to this maximum, it would larger than the numerical value shown in Table 6). This is because that if more investigation is not profitable, the open model will terminate exploration and make an immediate decision. Therefore, the open model is computational scalability that may not the case for the closed model.

Table 6. Average consultant fee of closed model vs. open model during the process of trust evaluation

Number of Agents	Average Consultant Fee		Number of Agents	Average Consultant Fee	
	Closed Model	Open Model		Closed Model	Open Model
100	1048.36	11.23	600	2044.53	13.23
200	1948.54	12.85	700	2034.93	13.08
300	2075.07	13.61	800	2267.41	13.38
400	2062.19	13.27	900	2287.80	13.57
500	2302.92	13.14	1000	2581.62	13.31

Varying the Accuracy of the Survey. It is necessary to know the effect that the accuracy of consultant agents' assessments of service quality has on the average precision of the system. We explored this by varying the variant of $\boldsymbol{X}$ (the accuracy of the survey) from 0.02 to 0.12. We set the number of agents to 100 here. As shown in Figure 5, we found that the more accurate the assessment, the more exactly we estimated value of trust. So, the closed and open trust models all depended on the accuracy of the survey. Also the high correlation between the precision of the open model and the accuracy of the survey was observed.

Varying the Population Quality. It is important to understand how the average precision is affected by the quality of agent population. We explored this by the varying the mean quality of agents and set the number of agents to 100 too. To measure the robustness of the models to bad agents, we selected agent qualities from six Gaussian distribution, with means from 0.4 to 0.9 and the same variant of 0.15. We varied the fraction of agents drawn from each distribution. Overall, as shown in Figure 6, we found that the system using the closed and open models differed only slightly in terms of precision, and the better the agent population, the higher the average precision was, which makes sense because in this case, the agent should get a more accurate estimate of provider's quality.

The results show that the closed and open trust models are robust to the quality of agent population and the accuracy of the survey generally, and the open trust model is robust to the number of agents. Also, the open model outperforms the closed model in terms of precision and consultant fee after the number of agents exceeds 700. We believe the reason underlies these results is: the closed model uses the transitivity strategy to prune its searches and is affected by the length of a chain of recommendations, falling as the chain gets longer. This pruning is likely to result in the loss of too much investigation. However, the idea of the open model is that at every stage of the procedure (after every given observation) one should compare the posterior Bayesian risk (or utility) of making an immediate decision with the "expected" posterior Bayesian risk (or utility) that will be obtained if more elaborate observations are taken. If it is cheaper to stop and make a decision, that is what should be done. Experimental results also show that the open model will stop mostly before pruning its all searches and its precision is not bad than the close model. Furthermore, the open model scales well to any social network size, as only tiny subsets of relatively constant size are visited and is computational scalability.

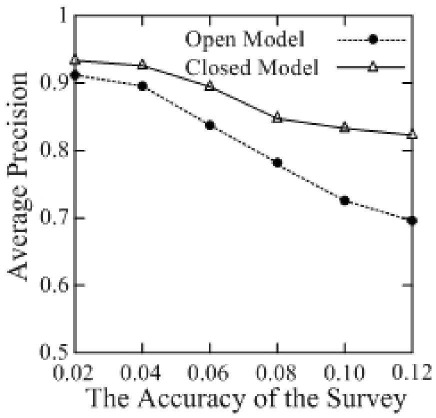

Fig. 5. Average precision for various accuracy of the survey

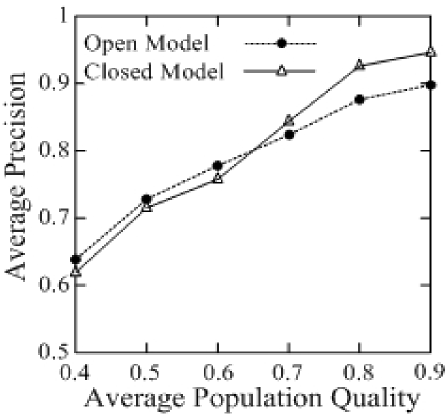

Fig. 6. Average precision for various qualities of agent population

5 Conclusions and Future Work

To achieve a pervasive, worldwide Semantic Web, an enhancement of computational trust model should be designed to support and enable computers and people to acquire, represent, exchange and integrate data and information available efficiently and conveniently. The Semantic Web is full of heterogeneous, dynamic and uncertain. If it is to succeed, trust will inevitably be an issue. After the cost and utility associated with the trust evaluation are discussed, two trust models have been formalized. We extend our trust models from discrete case of previous work to continuous case and improve the performance and precision of the system in this paper. Our work's contributions are: (1) The closed and open trust models have been proposed based on Bayesian sequential analysis. These models give trust a strict mathematical interpretation in terms of probability theory and lay the foundation for trust evaluation; (2) The utility and four types of costs: operational, opportunity service charge and consultant fee incurred during the process of trust evaluation have been considered sufficiently and an approach is proposed to balances these cost and utility; (3) Our approach enables users to combine a variety of sources of information to cope with the inherent uncertainties within the open Web environment and each user receives a personalized set of trusts, which may vary widely from person to person. (4) Experiments and some formal robust analyses have been made to examine the performance of our trust models. However, our proposed approach goes beyond other approaches in the kinds of representations of trust, the algorithms of trust evaluation and the formal analysis. Experimental results show that our open trust model is computational scalability and able to select the appropriate service providers for users effectively and efficiently according to their preferences.

For the future more robust analysis should be made properly, but we can mention that robustness can be dealt with more easily in Bayesian analysis. Service quality is multi-faceted. For instance, the file providers' capability can be presented in various aspects, such as the download speed, file quality and file type. We would like to consider multi-valued trust in the future.

Acknowledgment

The work is supported by the Subprogram of China 973 Project under Grant No. 2003CB317006, the National Science Fund for Distinguished Young Scholars of China NSF Program under Grant No. NSFC60533040, the Program for New Century Excellent Talents of China under Grant No. NCET-04-0545, and also the China NSF Program under Grant No. NSFC60503018.

References

1. B. Yu, M. P. Singh. Trust and reputation management in a small-world network. In Proc. 4th Int. MultiAgent Systems, 2000, pp.449-450.
2. Dong Huynh, Nicholas R. Jennings, Nigel R. Shadbolt. Developing an integrated trust and reputation model for open multi-agent systems. In Proc. 7th Int. Workshop on Trust in Agent Societies, 2004, pp.65-74.
3. D. J. Watts, S. H. Strogatz. Collective dynamics of "small-world" networks. *Nature*, June, 1998, 393: 440-442.
4. E. M. Maximilien, M. P. Singh. An ontology for Web service ratings and reputations. Workshop on Ontologies in Agent Systems, 2nd Int. Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Melbourne, Australia, 2003.
5. James O.Berger. Statistical decision theory and Bayesian analysis (second edition). Springer-Verlag, New York Inc, 1985.
6. J. Golbeck, Bijan Parsia, James Hendler. Trust networks on the Semantic Web. Lecture Notes in Computer Science, September, 2003, pp.238-249.
7. K. O'Hara, Harith Alani, Yannis Kalfoglou, Nigel Shadbolt. Trust strategies for the Semantic Web. Workshop on Trust, Security, and Reputation on the Semantic Web, 3rd Int. Semantic Web (ISWC), Hiroshima, Japan, 2004.
8. Li Ding, Lina Zhou, Timothy Finin. Trust based knowledge outsourcing for Semantic Web agents. In Proc. 2003 Int. Web Intelligence, October, 2003, pp.379-387.
9. Lik Mui, Mojdeh Mogtashemi, Ari Halberstadt. A computational model of trust and reputation. In Proc. 35th Int. System Sciences (HICSS), January 7-10, 2002, pp.2431-2439.
10. S. Majithia, A. S. Ali, O. F. Rana, D. W. Walker. Reputation-based semantic service discovery. 13th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, 2004, pp. 297-302.
11. M. Richardson, Rakesh Agrawal, Pedro Domingos. Trust management for the Semantic Web. In Proc. 2nd Int. Semantic Web, Sanibel Island, Florida, 2003.

12. M. Winslett, T. Yu, K. E. Seamons, A. Hess, J. Jacobson, R. Jarvis, B. Smith, L. Yu Negotiating trust on the web. *IEEE Internet Computing*, 2002, 6(6): 30-37.
13. N. R. Jennings, An agent-based approach for building complex software systems. *Communications of the ACM*, 2001, 44(4): 35-41.
14. S. D. Ramchurn, Nicholas R. Jennings, Carles Sierra, Lluis Godo. A computational trust model for multi-agent interactions based on confidence and reputation. Workshop on Deception, Fraud and Trust in Agent Societies, 2nd Int. Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS), Melbourne, Australia, July, 2003.
15. S. Milgram. The small world problem. *Psychology Today*, 1967.
16. S. P. Marsh. Formalising trust as a computational concept. Ph.D. dissertation, University of Stirling, 1994.
17. Xiaoqing Zheng, Huajun Cheng, Zhaohui Wu, Yu Zhang. A computational trust model for Semantic Web based on Bayesian decision theory. In Proc. 8th Int. Asia Pacific Web Conference, January 16-18, 2006, pp.745-750.
18. Yao Wang, Julita Vassileva. Bayesian network-based trust model. In Proc. 2003 Int. Web Intelligence, October 13-17, 2003, pp.372-378.
19. Y. Gil, V. Ratnakar. Trusting information sources one citizen at a time. In Proc. 1st Int. Semantic Web (ISWC), Sardinia, Italy, June 9-12, 2002, pp.162-176.