Multiagent System for Dynamic Web Services Selection

E. Michael Maximilien
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120
maxim@us.ibm.com

Munindar P. Singh
North Carolina State University
Department of Computer Science
Raleigh, NC 27695
singh@ncsu.edu

ABSTRACT

Service Oriented Architectures (SOAs) promise to enable the creation of business applications from independently developed and deployed *services*—roughly, software components that encapsulate and provide business functionality through standardized interfaces. A key advantage of SOAs is that they enable services to be dynamically selected and integrated at runtime, thus enabling system flexibility and adaptiveness—autonomic attributes that are key for modern business needs. However, current techniques provide no support for actually making rational selections, which are key to accomplishing autonomic behavior.

We develop a multiagent framework based on an ontology for QoS and a new model of trust. The ontology provides a basis for providers to advertise their offerings, for consumers to express their preferences, and for ratings of services to be gathered and shared. The ratings are essential, because they give an *empirical* basis for the selection of services. The ratings are quality-specific and are obtained via automatic monitoring or, if appropriate, user input.

The agents thus form an ecosystem in which they help each other. We empirically evaluate the resulting system via simulation. Our results show that the agents are able to dynamically adjust their trust assignments and thus continually select the best available services for their consumers' needs.

Keywords

Autonomic Computing, Multiagent Systems, Quality of Service, Reputation, Trust, Semantic Web Services, Service Selection

1. INTRODUCTION

The SOA vision is that, first, providers will offer several (potentially competing) services and, second, prospective users of services will dynamically choose the best offerings for their own purposes. For example, you might choose the best hotel booking service or the best bookseller, where you alone decide what is best for you. Likewise, you might choose the best component services (such as logging, backing up, and so on) to construct and deploy an application that meets your needs.

However, current approaches only partially address the SOA vision. They enable services to be described and listed in public registries (analogous to telephone directories). But they provide no means of selecting among multiple services that appear to perform the same function. In other words,

you are forced to make an ad hoc decision about which of the many hotel booking services or booksellers to use. Because tens of thousands of specialized and not widely known services are involved, a practical approach cannot merely pre-select a few famous companies such as Amazon.com, but must apply at a much larger scale. Further, when we consider not only business services but also component services, the choices can become quite subtle because the same provider may offer multiple alternatives.

The thesis of this paper is that service selection can be rationally carried out only on an *empirical* basis—that is, how a given service has behaved, not only how it was advertised. Given the large number of services, users must share information about their experiences—in effect, multiplying the benefit of their empirical evaluations by sharing them. Traditional, proprietary reputation systems (such as those maintained by eBay) and proprietary recommender systems (such as those maintained by Amazon.com) are not suitable for services. In particular, they do not allow a customizable schema in terms of the qualities of interest to different users, interject themselves into each transaction, and own the data that is gathered.

What is needed is a means to allow service consumers to share quality opinions, which presuppose an agreed upon set of QoS definitions. Using these quality opinions, service consumers can derive the reputation of service implementations on these qualities. By knowing its quality needs for an application, a service consumer can derive a *trust* value for each available service implementation. Thus, selecting the best service implementation simply corresponds to selecting the most trusted implementation. We can automate the service selection task with software agents acting on behalf of service consumers.

1.1 Contribution

An important characteristic of automatic selection by using trust in open environments, such as the Web, is that trust should be *self-adjusting*. That is, service implementations that behave incorrectly should (in essence) be purged from the system by virtue of not being selected. Poor service implementations should accumulate a low reputation. Conversely, when a once awry service implementation starts to behave correctly, we would like the agents to increasingly consider it for selection. This dynamic and self-adjusting consideration of trust for selection matches the goals of autonomic computing [9].

Self-adjusting trust. The autonomic characteristic of a multiagent system whereby the levels of trust between

the interacting parties are dynamically established and adjusted to reflect recent interactions.

We develop a multiagent framework that uses an ontology for QoS to support self-adjusting trust. The ontology gives a means for providers to advertise their offerings, for consumers to express their preferences, and for ratings of services to be gathered and shared. The ratings yield an *empirical* basis for the trust placed in different implementations. Moreover, the agents thus form an ecosystem in which they help each other identify the best implementations. Poorly performing implementations can be avoided. The converse challenge is to introduce new services or revive services that behaved poorly but are now functioning well again. To this end, this paper introduces what we term *explorer agents*. The explorer agents provide a means to monitor different service implementations, especially those that are new or currently out of favor. They allow consumers to select implementations that are predicted to perform well along the qualities of interest to them, even if there is inadequate positive experience with such implementations.

Our evaluation shows that the agents are able to dynamically adjust their trust assignments and thus continually select the best available services for their consumers' needs.

1.2 Organization

The remaining of the paper is as follows. Section 2 gives an overview of our trust model and how it is used to solve the service selection problem. Section 3 briefly discusses the technical framework including highlighting previous results. Section 4 gives a detail empirical evaluation showing the emergence of self-adjusting trust. Section 5 highlights various related work in the field and Section 6 concludes and gives some directions for future work.

2. BACKGROUND

We developed an agent-based approach for service selection that includes a flexible notion of trust based on reputation. The agents transparently attach to existing services and enable their dynamic selection. We introduce a comprehensive *ontology* (roughly, a taxonomy with some additional features) for qualities of services [11]. This ontology includes the well-known computing qualities such as throughput and latency, but provides hooks to include any application-specific or even idiosyncratic qualities that users may need, such as shipping delay. Our agent-based framework enables users to share information about any of the qualities. We developed algorithms by which user preferences regarding which qualities they consider more or less important can be applied (using the reputation data) to help each user select services that best meet his needs. The framework also has the ability to continually monitor services so that services that begin to perform poorly (relative to a particular user's preferences) are de-selected and those that begin to perform well are re-selected.

We now give an overview of how we model the service selection problem, our QoS-based trust model, and show how our model can be used as a solution to the service selection problem.

2.1 Service Selection

We model each Web service $s = (\iota, i)$ as a pair, where $\iota \in \Upsilon$ is the interface and $i \in I_\iota$ is an implementation of the service. Υ represents the set of all URIs and $I_\iota \subseteq \Upsilon$ is the set of all service implementations of interface ι. For each service s we associate an application domain $d \in \Delta$ where Δ is the set of all application domains. An example of an application domain is *Math* representing Web services for mathematical calculations; another domain is *Finance* with services such as loan and stock quote.

With each application domain d we associate qualities $Q \in \Phi_d$ representing nonfunctional attributes common to the Web services in the application domain d. Φ is the set of all qualities. For each quality Q and service s we let $\widehat{Q_s} = \{q_1, \ldots, q_n\}$ be the set of collected opinions, on quality Q. These opinions corresponds to n selection of service s.

We assume without loss of generality that for each selection of service s we obtain quality opinions (from the selection agent) for each quality in the domain d of s. We can now formulate the service selection problem.

DEFINITION 1 (SERVICE SELECTION). *Let P be the set of all providers with implementations for interface ι. Our problem is to select the service implementation $i \in I_\iota$ of service s from all service providers $p \in P$ such that:*

$$i = \arg \max_{i \in I_\iota} \{trust(i, \Phi_d)\} \tag{1}$$

Where $trust() : I_\iota \times \Phi \mapsto \mathcal{R}$ is a service trust function.

2.2 QoS Preferences and Advertisements

Each service consumer $c \in C$, where C is the set of all consumers, will have service quality needs that are specific to its application. For instance, the consumer of a *StockQuote* service used to give quick security quotes on financial Web sites has different quality needs than a consumer of this same quote service but within a brokerage application used to buy and sell securities. The latter's need for fast response time, high availability, and accuracy are critical to the brokerage application's success.

In order to accurately select services for consumers we first need a means to represent the consumer's needs for each quality exposed by the service.

DEFINITION 2 (CONSUMER QUALITY PREFERENCE). *A consumer's preference for a quality $Q \in \Phi$ is given by $\pi = (\pi_{min}, \pi_{pref}, \pi_{max})$, where π_{min} is the minimum value acceptable for the quality, π_{max} the maximum acceptable, and π_{pref} is the preferred value for quality Q all from the perspective of the consumer. We require $\pi_{min} \leq \pi_{pref} \leq \pi_{max}$.*

A consumer's preferences for a service s is a collection of quality preferences, one for each quality needed by the consumer.

Since service providers may offer different services with specific targeted consumers, we also need a means to represent the providers' advertisements for each quality in the set of qualities exposed by a service.

DEFINITION 3 (PROVIDER QUALITY ADVERTISEMENT). *A quality advertisement α is a proclamation by a provider for a particular quality $Q \in \Phi_d$. More specifically, we denote $\alpha = (\alpha_{min}, \alpha_{typical}, \alpha_{max})$, where α_{min} is the minimum advertised value for the quality, α_{max} the maximum, and $\alpha_{typical}$ is the typical value promised for the quality Q by the provider. We require $\alpha_{min} \leq \alpha_{typical} \leq \alpha_{max}$.*

A provider's service advertisements for a service implementation is a collection of quality advertisements, one for each quality applicable to the service.

2.3 Trust Model

To provide a solution to Equation 1 we need a *trust()* function that uses the collected quality values while taking into account the quality preferences of the consumer and the advertisements of the provider. The resulting trust value for a service implementation would enable us to rank different service implementations according to how well they meet the consumer's quality needs.

Let us assume that the collected quality opinions values are normally distributed with minimum, maximum, mean, and variance which are inferred from domain experts and attached to the QoS ontology. We can then assume that the collected values $\widehat{Q}$, for each quality Q, can be normalized as Z statistics values [20, p. 173]. This enables us to meaningfully compare qualities and aggregate them into a single value.

With this assumption, we start by defining an aggregation of the collected quality values for a quality Q that represents the general opinions of all agents that have selected the given service implementation.

DEFINITION 4 (SERVICE QUALITY REPUTATION). *We denote the reputation $R_Q^{(i)}$ of a service implementation, with respect to quality Q, as the aggregation of the quality opinions (i.e., quality values) for the service implementation i of service s over some time interval. Specifically,*

$$R_Q^{(i)} = \frac{1}{n} \sum_{k=1}^{n} q_k \delta^{-t(q_k)} \qquad (2)$$

where n is the number of collected quality values, $\widehat{Q}_i = \{q_k\}_{k=1}^{n}$ is the set of quality values collected from service agents as they selected the service implementation i, $\delta \in R$ is the quality Q's dampening factor, and $t() : \Phi \mapsto \mathcal{Z}^+$ is the time for which the quality value q was collected. $t(q) = 1$ for the most recent collected value and $t(q) > 1$ for all other values.

With $R_Q^{(i)}$ we have a representation of the general opinion on how well service implementation i performs for a quality Q. We now need a means to derive the trust value that a prospective consumer of service s should assign to each service implementation i of interface ι using the quality reputations for implementation i, the consumer's quality preferences, and the quality advertisements of provider p of implementation i.

Since the quality advertisements and preferences are defined as points on the quality line of Q, we can calculate the moment of these points with respect to the π_{pref} of the consumer preferences. In essence, the closer the advertised values and reputation are to the preferred value, the greater the degree of match (and of the resulting trust). Generally, Equation 3 shows the second moment of a vector $\vec{x} = \langle x_1, x_2, \ldots, x_n \rangle$ about some point a.

$$moment(\vec{x}, a) = \frac{1}{n-1} \sum_{i=1}^{n} (a - x_i)^2 \qquad (3)$$

We formulate the consumer's trust assignment for an implementation using Equation 3. However, since we want to match service implementations whose advertisements match the needs of a service consumer, we start by defining a matching operator between quality preferences and advertisements.

DEFINITION 5 (PREFERENCE MATCHING OPERATOR ▷). *For each $Q \in \Phi_d$ let $\alpha_Q = (\alpha_{min}, \alpha_{typical}, \alpha_{max})$ is the advertisement of provider p of service implementation i for quality Q and $\pi_Q = (\pi_{min}, \pi_{pref}, \pi_{max})$ be the consumer's preferences for quality Q.*

We define the preference matching operator ▷ for Q as:

$$\pi \triangleright \alpha = \begin{cases} (\pi_{max} \leq \alpha_{max}) \wedge (\pi_{pref} \geq \alpha_{min}) \wedge \\ (\pi_{pref} \leq \alpha_{max}) \quad \text{if } dir(Q) = \uparrow, \text{ and} \\ \\ (\pi_{min} \leq \varphi_{min}) \wedge (\pi_{pref} \leq \alpha_{min}) \wedge \\ (\pi_{pref} \geq \alpha_{max}) \quad \text{when } dir(Q) = \downarrow \end{cases}$$

Where $dir(Q) : \Phi \mapsto \{\uparrow, \downarrow\}$ is associated with each quality Q, such that $dir(Q) = \uparrow$ indicates that the quality Q is directionally increasing which means that higher values for Q are generally preferred by service consumers. And $dir(Q) = \downarrow$ indicates that lower values are generally preferred.

Using the ▷ operator we can derive the *trust* function of Definition 1 as follows.

DEFINITION 6 (SERVICE TRUST FUNCTION). *Let $Q_{min} = \min(\alpha_{min}, \pi_{min})$ and $Q_{max} = \max(\alpha_{max}, \pi_{max})$.*
Let $\vec{Q}_i = \langle Q_{min}, \alpha_{typical}, \pi_{pref}, Q_{max}, R_Q^{(i)} \rangle$.

$$\begin{aligned} qTrust(\vec{Q}_i, \pi_{pref}) &= moment(\vec{Q}_i, \pi_{pref})^{-\frac{1}{2}} \\ &\quad where \ moment(\vec{Q}_i, \pi_{pref}) \neq 0 \\ serviceTrust(i) &= \sum_{\substack{Q \in \Phi_d, \\ \pi \triangleright_Q \alpha}} qTrust(Q_i, \pi_{pref}) \\ trust(i_p, \Phi_d) &= serviceTrust(i_p) \end{aligned}$$

Where $i_p \in I_\iota$ is an implementation of service s by provider p.

3. FRAMEWORK

To evaluate our trust model and hypothesis of self-adjusting trust, we created a framework that augments a typical SOA with agents. The principal idea is to install software agents between service consumers and each service that they consume. These consumer service agents expose the same interface as the service. However, they augment the service interface with agent-specific methods. An example of such a method is *setWsPolicy()* which allows consumers to communicate their QoS preferences. By exposing the same interface as the service these agents are able to transparently and dynamically select the actual service implementation by considering the service consumer's quality needs. The consumer communicates its needs via the augmented agent interface. Service method invocations are done via the service agent who in turn monitors and forwards all calls to the selected service.

A high-level view of the architecture; the details and runtime operation are described at length in [11]. Briefly, an SOA application makes use of some service which has many implementations by different providers. Instead of selecting the implementation directly, the application uses a service

agent which exposes the same interface as the service and selects, on the consumer's behalf, the implementation which best matches the consumer's policy.

In addition, the service agents participate in common *agencies* where they share their quality opinions on the selected service implementations. An agency is simply a *rendezvous* node on the network where quality opinions are shared and aggregated.

The agents share a conceptualization of quality in the form of an ontology. The ontology is partitioned into three parts. The *upper QoS ontology* contains basic definitions for all qualities, including modeling relationships between qualities. The *middle QoS ontology* extends the upper ontology and defines qualities that are applicable across different domains. *Lower QoS ontologies* are defined for specific domains by extending qualities in the middle ontology or creating new ones from the upper ontology. Service agent behaviors for quality monitoring can also be attached to the ontology and dynamically bootstrapped in the agents. Maximilien and Singh [11] give an overview of the upper and middle QoS ontology as well as discussing examples of lower ontology qualities and example usages of the framework.

We implemented this architecture in the Web Service Agent Framework (WSAF) and used simulation experiments on simple services as an initial evaluation [12]. The initial results showed that the service agents are able to accurately select service implementations according to the consumer's preferences and adjust the selection as service implementations' quality degrade. In Section 4, we evaluate an approach to enable the service agents to adjust their service selection when a well-behaved service implementation starts degrading its exposed qualities and then again provides good qualities (or vice versa).

4. EMPIRICAL EVALUATION

Our empirical evaluation consists of a series of simulations tailored to show the emergence of self-adjusting trust. The simulations results reveal how, empirically, our trust model yields a system that autonomically adjusts the level of trust for the service implementations depending on their past quality behaviors. Previous experiments showed that as service qualities drop, the consumers' agents select other service implementations and eventually converge to the 'best' service implementations [12]. In the current experiment we expand on the previous one by adding *explorer agents*. The simulations reveal the importance of the explorer agents in the agent community and also how recent quality information can be captured and be benefited from.

4.1 Setup Summary

Briefly, we created three sets of a simple mathematical integer sorting service with one method to sort an array of integers. Each set contains five identical service implementations. Each implementation exposes the qualities: *MethodFaultRate*, *PercentAvailability*, and *MethodResponseTime*. The qualities represent the average fault-rate of service methods, the average service availability (as a percentage), and the average response-time of the service's method. In addition the implementations of each set have identical quality advertisements. We created three groups of five consumers; each group has quality preferences biasing the members to one of the set of service implementations. We named the groups of consumers: *Careful*, *Mellow*, and *Rushed*; indicat-ing their general preferences biased for the three qualities.

We ran various simulations attaching a service agent to each consumer and collected the agent's selections. The service implementations in each group are numbered $0, \ldots, 4$ as are the service consumers. In some of the simulations, we artificially forced all service implementations of a group, except the last one (numbered 4), to have its quality degrade. We term this *doping* the implementations. Our results show that the service agent is able to find and select the clean service implementation in a group after some number of iterations to build the necessary reputation of service implementations.

For this experiment we run simulations with a similar setup of service implementations and consumer service agents as in the previous experiment; however, for each simulation we:

- Added *explorer agents*. These are consumer service agents deployed for a specific service interface that periodically run an exploring task on the service implementations. These agents' primary purpose is to explore the community of services regardless of a service's trustworthiness. That is, these agents would select a service implementation whose quality reputation has dropped to a point where it is no longer selected by regular consumer agents, even though more recently its quality characteristics may have improved. The explorer agents do not discriminate and have no preset quality preferences. They select the available service implementations in a round-robin fashion. For each implementation, they execute the same exploring task. They monitor the selected service for all of its advertised qualities and participate in appropriate agencies for these qualities. The net effect is that if a service implementation quality has improved, the explorer agents will help it gain positive reputation. The primary controlling factor for an explorer agent is how frequently it performs its exploring tasks. We name this parameter the *execution frequency* of the explorer agent. Another controlling factor is the number of explorer agents as a ratio of the total number of agents in the community.

- Varied the service *doping policy*. As in the previous simulations we artificially control the qualities of certain service implementations with a quality doping mechanism. This allows us to have knowledge (as an oracle) to which service implementation should be selected by a service agent at any given time during the simulation. To show the effect of the explorer agents, we also control the duration of the doping. We introduce two new doping parameters: *Stop doping after* and *Restart doping after*. They specify, as their names suggest, when doping of the service is to stop and to be restarted—enabling service implementations to have periods of correct and incorrect service qualities.

To show the emergence of self-adjusting trust, we conduct seven simulations, varying the number of explorer agents and their execution frequency. We also simultaneously vary the above doping parameters to enable us to better predict and measure the effects of the explorer agents.

Tables 1 shows an overview of the simulations with regard to the explorer agents and related parameters. The

columns show the simulation number, the number of iterations, whether the service implementations are doped or not, the number of explorer agents, the explorer agents ratio to the total number of agents in the simulation, and the explorer agents' execution frequency.

Table 1: Simulations' explorer agents parameters.

Sim.	Iter.	Is doped?	No.	Ratio	Exe freq. (ms)
1.0, 1.1	15	No	0	0.0	NA
1.2	15	Yes	3	0.2	4500
1.3	10	Yes	3	0.2	6000
1.4	10	Yes	3	0.2	5000
1.5, 1.6	10	Yes	3	0.2	4500

Table 2: Simulations' doping policy parameters.

Sim.	Doping	Stop doping	Restart doping
1.0	NA	NA	NA
1.1, 1.2	Full doping	Never	Never
1.3, 1.4, 1.5	Full doping	After 6	Never
1.6	Full doping	After 6	After 15

Table 2 shows the doping policy variation for each simulation. The *Doping* column indicates the overall service doping policy used. *Full doping* refers to the same service doping policy as in the previous experiment [12] and with changes summarized in Table 3. The *Stop doping* column indicates the number of service selections after which service doping is stopped. The *Restart doping* column indicates the number of service selections after which service doping is resumed after being stopped.

Table 3: Full-doping. Doped quality names and any other additional doping parameters. Doping affects all service implementations except for the last one in each set.

| Doped quality name	Parameters
FaultRate	Random periodic, $p = 3$
PercentAvail	Random periodic, $p = 4$
InvokeTime	Random delay, $10\ ms < d < 20\ ms$

4.2 Results

We now further describe each simulation and its expected results. For each simulation we also show the results obtained in the form of service selection graphs. The simulations are designed to progressively show the impact of the explorer agents on the system as well as the emergence of self-adjusting trust.

Sim 1.0, 1.1, and 1.2: Base line. These simulations establish the base line service selection with and without explorer agents. As in the previous experiment we expect each group of consumers to eventually converge to the sole clean service implementation [12]. Our expected results for both simulations, with and without explorer agents, should be similar to Figure 1 (from previous experiment).
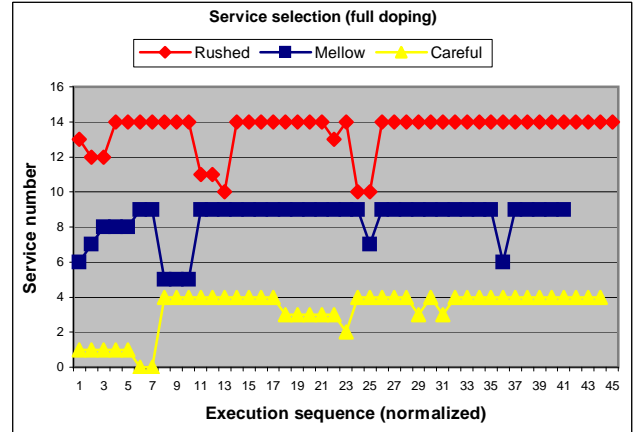
Figure 1: Selection with full doping.
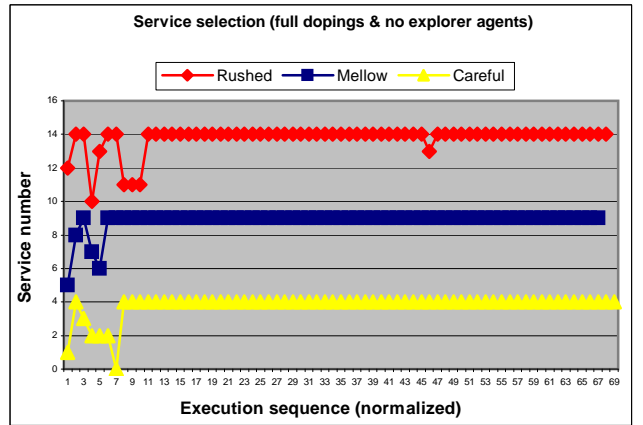
Figure 2: Full service doping and no explorer agents (Simulation 1.0).
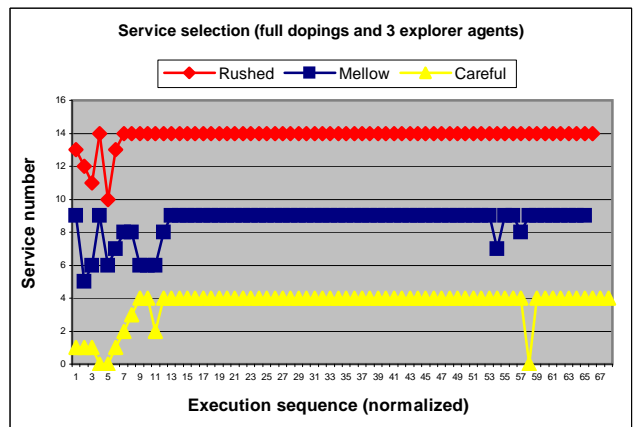
Figure 3: Full doping and three explorer agents (Simulation 1.2).

Figures 2 and 3 show the results obtained for simulations 1.0 and 1.2. Simulation 1.1, where doping is stopped after six selections has analogous results. These are in agreement with our expectations. The consumer service agents, after a period of trying all service implementations in the same pool, eventually find the sole clean service implementation and converge to it. This comes about when no explorer agents are present and the doping lasts the entire simulation (Figure 2) and also when the doping stops. This occurs because the consumer agents never select any of the newly well-behaved service implementations—which accumulated a low reputation already and are never considered.

In Figure 3, even though there are explorer agents present, since the doping never stops, the sole clean service implementation remains the best alternative for selection during the entire simulation.

Sim 1.3, 1.4, and 1.5: Explorer agents. There are two principal goals for simulations 1.3 to 1.5. First, show that when we stop doping the service implementations pool, the use of explorer agents (EAgents) will eventually cause the consumer agents to reselect services that were avoided when they were doped. Second, get some understanding of the impact of the explorer agents' execution frequency on the reselection of these newly cleaned service implementations.

Figures 4 and 5 show the service selection results with three explorer agents and execution frequency reduced progressively from 6000, and 4500 ms (the result for 5000 ms is similar). The first thing to notice is that in each case, the consumer agents first converge to the sole clean service implementation (as in simulations 1.0 and 1.1). Then as the stop doping takes effect—the service implementations start to have good qualities—the effect of the explorer agents takes hold and the service implementations start to be selected randomly again. This is the emergence of self-adjusting trust—as these service implementations behave correctly (they are not doped) they are reconsidered for selection.

We also notice that with higher execution frequency, after the service doping is stopped, the newly well-behaved service implementations are more often selected than with lower execution frequency. This can be noticed by comparing Figure 5 with Figure 4 in the part of the graph where reselection is occurring. We conjecture that this is due to faster positive reputation accumulation for the well-behaved service implementations due to frequent explorer agent selections.

Sim 1.6: EAgents with doping restart policy. In this simulation, we use the best setup from Simulations 1.3 to 1.5, but in addition we use a doping restart policy. That is, we simulate service implementations behaving incorrectly again. The restart is set up to occur after the 15th selection. We expect the consumer agents to begin again converging to the sole clean service implementation.

Figure 6 shows the results of this simulation. As expected, the consumer agents first converge to the sole service implementation. Then as the doping stops


Figure 4: Full service doping stopped after six and three explorer agents with 6000 ms execution frequency (Simulation 1.3).
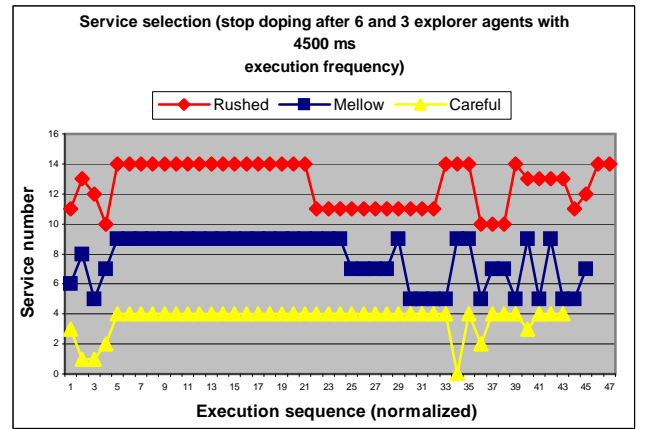

Figure 5: Full service doping stopped after six and three explorer agents with 4500 ms execution frequency (Simulation 1.5).

they start selecting other service implementations—since the explorer agents help these service implementations rebuild their quality reputation and thus their trustworthiness. As the services doping restart, the consumer agents again gradually converge to the sole clean service implementation. This result is in agreement with the emergence of self-adjusting trust.

5. RELATED WORK

We divide the works related to ours into two broad categories: QoS-based service selection and trust and trustworthiness. For each category we discuss the ones that are closest to ours, highlighting similarities and differences.

5.1 QoS-Based Service Selection

Service selection approaches fall into two primary categories: design-time and runtime. In design-time selection, the application designer or architect use service registries coupled with service descriptions to select and test binding to a service. Nonfunctional characteristics are considered

during trial and error tests of the selected services. Newer techniques using richer semantic descriptions of services can help in the discovery of service interfaces. OWL-S [16] is an example of a rich service ontology used for semantic service discovery.

In runtime service selection, which is of greater relevance to our work, the service interface is already discovered, but the service implementations must be discovered, selected, and bound to, based on nonfunctional service attributes. Work in this area encompasses QoS requirements, models and metrics, and middleware. The W3C QoS Requirements for Web Services [19] gives an overview of the requirements. The QoS UML profiles described in [1] and [3] are examples of QoS models. Various brokering and middleware architectures [15, 13, 14, 17, 21, 24] have been proposed for using QoS in Web services.

Other researchers have also proposed using QoS for service selection [13, 17, 7]. However, whereas these works address some form of service selection, they do not address it adequately for open environments, such as using trust and reputation and a decentralized multiagent architecture as we are proposing. Wohlstadter et al. [21] propose a policy language for advertising the QoS needs of clients and to allow the middleware to match servers and clients. However, their work lacks a complete conceptualization of nonfunctional attributes for Web services. Further, their matchmaking techniques is not geared to enable dynamic evolution of the QoS exposed by the services. In other words, it does not exhibit autonomic characteristics.

Casati et al. [2] present a dynamic service selection using a data mining approach on the service conversation logs. They assume that conversation logs are collected for each business process execution and that service consumers specify quality goals that can be measured from the logs. Using data mining they provide a middleware that can dynamically analyze the logs of various conversations and determine the services best matching the service consumer's goals. Though similar in intent, our approach differs in two respects. First, in our architecture, we use a multiagent system with an agreed upon QoS ontology, whereas Casati et al. use a centralized middleware and their definitions of qualities is derived from conversation logs. Second, our selection is based on the trust value assigned to a service implementation, which itself is based of a combination of the reputation of the consumer's quality needs. Casati et al.'s selection is based on a service ranking which is derived from the decision trees constructed from the conversation logs.

5.2 Trust and Trustworthiness

The literature on trust in the context of Web systems and Web services has been growing lately. The literature closest to our work is summarized in [5]. Though a bit dated, most of the techniques described are applicable to our framework. It is well-known that no security techniques can apply in a purely open environment where identities can be readily created and discarded. In such cases, the best we can do is to use social approaches based on open reputation management. Yu and Singh have examined the robustness of such approaches under certain threats [22].

Other relevant trust approaches are the Pretty Good Privacy (PGP) Web of Trust and the Platform for Internet Content Selection (PICS) [18]. However, these works are mainly in the realm of creating trust systems and platforms
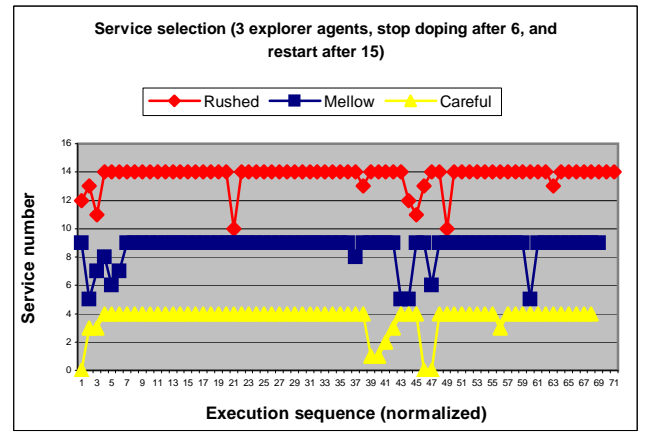

Figure 6: Three explorer agents with full service doping stopped after six selections and restarted after 15 (Simulation 1.6).

for human users and human-facing applications. Our goal is to create trust where human involvement is limited or completely absent.

Work on reputation systems is also closely related to our work. Our proposed reputation calculation extends that of Zacharia and Maes [23] and applies it in the world of Web services by making considerations of QoS explicit. Huynh et al. [6] calculate interaction trust as a reputation value but their approach is limited when applied to service selection (they do not consider nonfunctional characteristics of services) and they also have limited prospect to creating a system that is autonomic. Giorgini et al. [4] present an agent-based social and individual trust model. Their model is mainly in the realm of requirements engineering and gathering; however, the resulting model is a step in the right direction to achieving a formal agent-based trust model and could be used to complementent our own. Kalepu et al. [7] extend our own approach to consider variance in collected ratings values and it is also specific to reputation for qualities that are subjectively rated.

Kashyap [8] argues that trust models for information systems and service composition must have a verification process. He gives a partial taxonomy of trust that includes transitive trust (endorsements), reputation (indirect trust), and *a priori* beliefs. Kashyap sketches how a composed trust quality value can be computed from the various trust notions. We differ primarily with this work in that we compute an aggregate trust value from QoS but agree with the fact that a complete model for trust must be multidimensional.

6. CONCLUSIONS

The autonomic characteristic of *self-adjusting trust* is crucial in open systems. We showed how explorer agents enable our trust model and multiagent system to exhibit this characteristic. In this manner, we can enable the benefits of autonomic computing for a wide range of settings.

Further aspects of the emergence of self-adjusting trust remain to be studied. These include, for instance, understanding the density of explorer agents needed to achieve a certain level of self-adjusting trust in the presence of ser-

vices advertising some qualities. Since explorer agents consume precious resources, an analytical model of their behavior in the community with respect to the resulting perceived level of self-adjusting trust would be valuable. Such a model would depend upon a characterization of the variability in practice in the behavior of service implementations.

The full exploitation of explorer agents may require the emergence of new standards. Specifically, explorer agents would not automatically be able to use services that require a consumer's private data or require payments. Agents that can perform such tests would be difficult to set up and specialized agencies may emerge to evaluate expensive services on the behalf of prospective consumers. However, if such services expose an *exploration interface* that enables exploration and evaluation of their offerings but without requiring payment, they would facilitate explorer agents. An exploration interface would be useful to a provider, because it would yield greater exposure to their offerings. An associated aspect is that the consumers will have to trust that the behavior obtained during exploration is representative of the behavior to be expected during the actual interaction.

Finally, a direction for our work is to improve the trust model to take into account provider's trustworthiness. Providers can be trusted by consumers *a priori* using implicit trust as well as indirect and transitive trust such as endorsements from trusted third parties [10].

7. REFERENCES

[1] J. Ø. Aagedal, M. A. de Miguel, E. Fafournoux, M. S. Lund, and K. Stolen. UML Profile for Modeling Quality of Service and Fault Tolerance Characteristics and Mechanisms. TR 2004-06-01, OMG, June 2004.

[2] F. Casati, M. Castellanos, U. Dayal, and M.-C. Shan. Probabilistic, Contex-Sensitive, and Goal-Oriented Service Selection. In *Proc. of 2nd Intl Conf. on Service Oriented Computing (ICSOC)*, pp. 316–321, Nov. 2004.

[3] V. Cortellessa and A. Pompei. Towards a UML Profile for QoS: A Contribution in the Reliability Domain. In *Proc. of the 4th Intl Workshop on Software and Performance*, pp. 197–206, 2004.

[4] P. Giorgini, F. Massacci, J. Mylopoulos, and N. Zannone. Modeling Social and Individual Trust in Requirements Engineering Methodologies. In *Proc. 3rd Intl Conf. on Trust Management*, Lecture Notes on Computer Science. Springer Verlag, 2005.

[5] T. Grandison and M. Sloman. A Survey of Trust in Internet Application. *IEEE Communications Surveys & Tutorials*, 3(4):2–16, 2000.

[6] D. Huynh, N. R. Jennings, and N. R. Shadbolt. Developing an Integrated Trust and Reputation Model for Open Multi-Agent Systems. In *Proc. of 7th Intl Workshop on Trust in Agent Societies*, pp. 66–74, July 2004.

[7] S. Kalepu, S. Krishnaswamy, and S. W. Loke. Verity: A QoS Metric for Selecting Web Services and Providers. In *Proc. of the 4th Intl Conf. on Web Information Systems Engineering Workshops (WISEW'03)*, pp. 131–139, Dec. 2003.

[8] V. Kashyap. Trust, But Verify: Emergence, Trust, and Quality in Intelligent Systems. *IEEE Intelligent Systems*, 19(5):85–87, Sept. 2004.

[9] J. O. Kephart and D. M. Chess. The Vision of Autonomic Computing. *IEEE Computer*, 36(1):41–50, Jan. 2003.

[10] E. M. Maximilien and M. P. Singh. Reputation and Endorsement for Web Services. *SIGecom Exchanges*, 3(1):24–31, Dec. 2001.

[11] E. M. Maximilien and M. P. Singh. A Framework and Ontology for Dynamic Web Services Selection. *IEEE Internet Computing*, 8(5):84–93, Sept. 2004.

[12] E. M. Maximilien and M. P. Singh. Toward Autonomic Web Services Trust and Selection. In *Proc. of 2nd Intl Conf. on Service Oriented Computing (ICSOC)*, pp. 212–221, Nov. 2004.

[13] S. Ran. A Framework for Discovering Web Services with Desired Quality of Service Attributes. In L.-J. Zhang, editor, *Proc. of the Intl Conf. on Web Services*, pp. 208–213, Las Vegas, June 2003

[14] S. Ran. A Model for Web Services Discovery with QoS. *SIGecom Exchanges*, 4(1):1–10, 2003.

[15] A. Sheth, J. Cardoso, J. Miller, and K. Kochut. QoS for Service-Oriented Middleware. In *Proc. of the 6th World Multiconference on Sytemics, Cybernetics, and Informatics (SCI02)*, vol. 8, pp. 528–534, July 2002.

[16] K. Sycara, M. Paolucci, A. Ankolekar, and N. Srinivasan. Automated Discovery, Interaction, and Composition of Semantic Web Services. *Journal on Web Semantics*, 1(1):27–46, Sept. 2003.

[17] M. Tian, A. Gramm, T. Naumowicz, H. Ritter, and J. Schiller. A Concept for QoS Integration in Web Services. In *Proc. of the 4th Intl Conf. on Web Information Systems Engineering Workshops (WISEW'03)*, pp. 149–155, Dec. 2003.

[18] W3C. Platform for Internet Content Selection (PICS), 1998. Recommendation.

[19] W3C. QoS for Web Services: Requirements and Possible Approaches, Nov. 2003. Note.

[20] D. D. Wackerly, W. M. III, and R. L. Scheaffer. *Mathematical Statistics with Applications*. Duxbury, Pacific Grove, 6th ed., 2002.

[21] E. Wohlstadter, S. Tai, T. Mikalsen, I. Rouvellou, and P. Devanbu. GlueQoS: Middleware to Sweeten Quality-of-Service Policy Interactions. In *Proc. of 26th Intl Conf. on Software Engineering (ICSE 2004)*, pp. 189–199, Edinburgh, May 2004.

[22] B. Yu and M. P. Singh. An Evidential Model of Distributed Reputation Management. In *Proc. of Autonomous Agents and Multi Agent Systems (AAMAS)*, pp. 294–301, 2002.

[23] G. Zacharia and P. Maes. Trust Management Through Reputation Mechanisms. *Applied Artificial Intelligence*, 14:881–907, 2000.

[24] L. Zeng, B. Benatallah, A. H. H. Ngu, M. Dumas, J. Kalagnanam, and H. Chang. QoS-Aware Middleware for Web Services Composition. *IEEE Transactions on Software Engineering*, 30(5):311–327, May 2004.