

Self-Adjusting Trust and Selection for Web Services

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 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 agents form an ecosystem in which they help each other select services best matching their preferences. We evaluate by simulation the agents' effectiveness in maintaining self-adjusting trust.

1. Introduction

Service-oriented architectures (SOAs) offer an essential ingredient of autonomic computing (AC). This is especially so when AC is understood generally as producing continual automatic adjustment to all aspects of a system, including its software configuration. To our way of thinking, the power of AC lies not just in managing data centers but more broadly in computing, including for business services. SOAs enable the flexible binding of services—i.e., well-encapsulated modules of functionality—to construct applications and systems that best meet user needs.

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. We developed a framework in which consumers can share quality opinions, thus creating a reputation for each service with respect to applicable qualities. Further, based on its local needs, a consumer computes a *trust* value for each available implementation. The best implementation for a consumer is

the one that it trusts the most. This work builds on our previous works: (1) an extensible ontology to model QoS [1], on which consumers' preferences and providers' policies are based; and (2) a mathematical formulation of trust based on consumers' preferences and providers' advertisements [3].

To support automatic selection in open environments, trust should be *self-adjusting*. That is, implementations that behave incorrectly should be purged by virtue of not being selected. Poor implementations should acquire a low reputation. Conversely, when a once awry service implementation starts to behave correctly, we would like to consider it for selection.

Self-adjusting trust. The autonomic characteristic of a multiagent system whereby the levels of trust between the interacting parties are dynamically established and continually adjusted.

2. Empirical Evaluation

Previous experiments involving our model showed that, as service qualities drop, the consumers' agents select other service implementations and eventually converge to clean implementations [2]. The current experiment introduces *explorer agents* and is designed to understand how they engender self-adjusting trust. Explorer agents periodically run an “exploring” task on the implementations of a given interface, thus checking their quality. The idea is that implementations that are not sufficiently trusted would still be explored, thus getting an opportunity to become trusted.

2.1. Experimental Setup Summarized

The simulations attached a service agent to each consumer and collected the agent's decisions. We created three groups of providers. We call artificially degrading a service implementation *doping* it. In some simulations, we doped all but one implementations of a group (the clean implementation shows as the top-most one for each line in the graphs below).

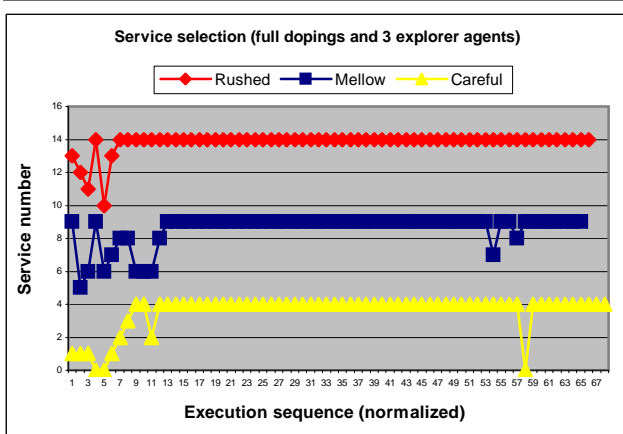


Figure 1. Full doping and three EAgents.

Our results show that, after some iterations, the consumer’s service agent finds and selects the sole clean implementation of a desired group. We modify the previous experiment by: (1) adding explorer agents; and (2) varying the service *doping policy*. We introduce two doping parameters: (1) *stop doping after*, which specifies the number of times a service is to be doped; and (2) *restart doping after*, which specifies when doping is resumed.

2.2. Results

We vary the doping policy of service implementations while retaining a clean implementation in each group.

Simulation 1.0, 1.1, and 1.2: Base line service selection.

These simulations establish the base line service selection with and without explorer agents. As in the previous experiment [2], we expect each group of consumers to eventually converge to the sole clean service implementation.

In Figure 1, even though explorer agents are present, since the doping never stops, the sole clean service implementation remains the best alternative for selection during the entire simulation.

Simulation 1.3: Explorer agents (EAgents). The main goal for this experiment is to show that when we stop doping the service implementations pool, the use of explorer agents will eventually cause the consumer agents to reselect services that were avoided when they were doped. A secondary goal is to see show that when the pool is doped anew, convergence to the sole clean service reoccurs.

Figure 2 shows the service selection result with three explorer agents. Notice that, the consumer agents first converges to the sole clean service implementation. Then as the doping stops—the implementations

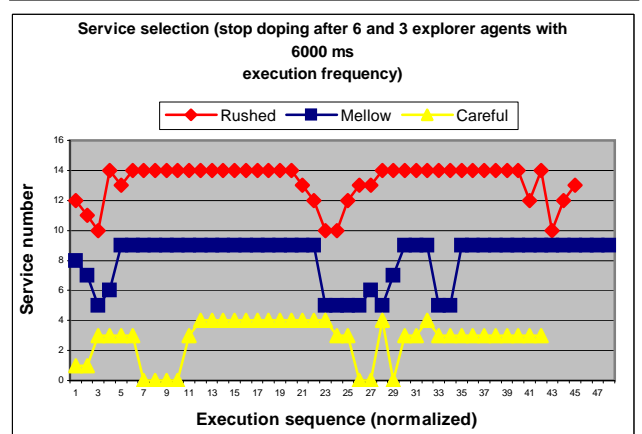


Figure 2. Doping stopped and restarted.

start to have good quality—the explorer agents kick in and the implementations start to be selected again. As the doping restarts then convergence to the clean service occurs anew. This is the emergence of self-adjusting trust.

3. Conclusions

Further aspects 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 services advertising some qualities. Since explorer agents consume precious resources, an analytical model of their behavior with respect to the resulting perceived level of self-adjusting trust would be highly 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. An exploration interface would be useful to a provider, because it would yield greater exposure to their offerings.

References

- [1] E. M. Maximilien and M. P. Singh. A Framework and Ontology for Dynamic Web Services Selection. *IEEE Internet Computing*, 8(5):84–93, 2004.
- [2] E. M. Maximilien and M. P. Singh. Toward Autonomic Web Services Trust and Selection. In *2nd Intl Conf on Service-Oriented Computing*, pp. 212–221, New York, 2004.
- [3] E. M. Maximilien and M. P. Singh. Agent-Based Trust Model Using Multiple Qualities. In *4th Intl Autonomous Agents and Multi Agent Systems*, Utrecht, Holland, 2005.