Ensuring Fair Trust in Reputation Based Systems

Radu Sion
Computer Sciences, Purdue University

In this paper we propose improvements to existing reputation based trust paradigms in distributed environments such as peer to peer networks [2] [3] and grids [1].

More specifically, we deal with two important aspects of such systems, namely (i) enabling the ability of a certain party to challenge a given rating by another party and (ii) ensuring a gradual decline in importance of “old” ratings, allowing for “comebacks” of “banned” parties.

Motivation. Enabling ratings-based trust management in distributed computing economies, in which parties are paid for services, brings about interesting questions, including: Are there any incentives for competing parties to rate each-other truthfully for service interactions? We would argue that indeed, in the traditional scenario, these incentives do not exist. Moreover, due to the competitive nature of these systems, malicious parties would actually benefit from offering negative ratings if asked about the competition. This can effectively result in denial of service attacks (to the rated parties) and/or disrupted flow of business (e.g. payed requests for computing cycles). We can differentiate two cases here: (i) if any and all parties can provide the same services (or have a significant overlap in their abilities, thus they are effectively competing for associated “customers”) and (ii) if, for a given service, there is only a limited subset of the parties that can perform it.

Malicious Competition. In the first case, because of the counter-incentive (to give correct ratings) illustrated above, mechanisms need to be envisioned in which for each transaction between server party S and client C for which is going to issue a rating, there exists a third trusted party (or possibly composed of $k + 1$ other parties, if we adopt a threshold-based assumption that no more than k of them are “bad”) which somehow “witnesses” the rating and can vouch for it (e.g. by a signature, similar ideas are explored in [3]). Lets use W to denote this witness. If W somehow observes the service interaction between S and C then it could possibly sign something like ”I, W saw that S delivered a certain result to C within the promised time-frame. Thus the rating for this interaction should be positive”. This raises other interesting question, such as: can W vouch also for the actual “accuracy” of the result, or other arbitrary “quality of service” features? If not, is the mere external “witnessing” worth anything? For W to be aware of the result accuracy, does it need to be able to actually compute it in parallel?

Multi-level Ratings and Clustering. Fortunately, in the second case, where not all parties can provide the same services, a certain solution can be deployed that solves some of the above issues. In this solution we propose to cluster the existing parties based on the ability to provide a certain service and then de-compose the rating scheme to the service level.

This makes sense because trust is not an all-or-nothing proposition. This is especially true in distributed environments (e.g. grids and peer to peer systems) where the same parties can exhibit different roles and perform multiple tasks (e.g. offer multiple services). Trusting a certain party to perform linear regression analysis correctly should not necessary imply it can also perform association rule mining (even if it is advertised) or that any of these services are performed in a timely manner.

Then, a party C should be able to only rate non-competing parties, that is parties that do not offer any of the services C offers. Of course a certain circularity still arises here an parties could collude and influence each other’s ratings by promises of other better services in the future etc.

Rating Decay and Challenges. Another issue of significant interest here is ensuring the “comeback” ability of once un-trusted or un-reliable parties that underwent a change. Because it is (arguably) un-reasonable to assume identity changes, a certain party can be “stuck” with a negative rating for a long time. If this rating also results in a lack of further interaction, e.g. a grid scheduler never decides to schedule jobs on a certain back-end CPU again, a closed “isolation” loop is complete. It might be of benefit to weigh ratings based on their “age” (i.e. when their associated transaction occurred) such that older ratings count increasingly less. This would “break” the isolation loop and allow the party to redeem itself.

Additionally we propose the introduction of an explicit redemption mechanism in which the given party (aware of its bad rating) broadcasts a claim message in which it challenges the rating and invites the rating parties to use its services and issue new associated ratings. This new claim could of course be false, however (in both cases) this issue can be solved by introducing an exponential back-off delay mechanism (in which exponentially-increasing delays are introduced after every deceiving claim) would curb truly malicious parties fast.

References

[1] Beulah Alunkal, Ivana Veljkovic, Gregor von Laszewski, and Kaizar Aminand. Reputation-based Grid Resource Selection. In *Proceedings of the Workshop on Adaptive Grid Middleware*, New Orleans, LA, September 2003.

[2] Ernesto Damiani, De Capitani di Vimercati, Stefano Paraboschi, Pierangela Samarati, and Fabio Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *Proceedings of the 9th ACM conference on Computer and communications security*, pages 207–216. ACM Press, 2002.

[3] Sepandar D. Kamvar, Mario T. Schlosser, and Hector Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *WWW*, 2003.