A Behavioural Model for Consumer Reputation

Anirban Basu, Ian Wakeman and Dan Chalmers

Department of Informatics, University of Sussex
{A.Basu, ianw and D.Chalmers}@sussex.ac.uk

In recent years, substantial academic research has focussed on solving the problem of email spam using reputation mechanisms (e.g., [1] and [2]) while others [3] have discussed the generalised problem of any unsolicited network transaction in client-server scenarios. We propose a reputation model based on behavioural history as a solution to this problem.

We use behavioural history of long-lived network identities to develop local and global reputation. This helps the provider to make decisions regarding future network transactions. The decisions are either binary (e.g., provide service or cut-off) or more fine grained where the provider can vary levels of service.

A declarative specification of behaviour quality of network identities in relation to relevant terms in the service contract is being developed, which will relate network activities to discrete units of good or bad behaviour. In most cases, clients have their own identities or they are identified as members of a group (e.g., customers of a particular Internet Service Provider). In situations where the identity of the client is anonymous and cannot be associated with any group, best effort service will be provided as a fall-back option. The term *score* is used to denote reputation value of the consumer, while the term *rank* is used for the provider. We have experimented with mathematical functions to best represent the reputation response to behaviour. We are using exponential saturation, exponential decay and linear patterns with adjustable slope (change) parameters. A local reputation response (ordinate) in relation to cumulative behaviour (abscissa) is illustrated in Fig. 1. The score saturates at positive or negative saturation levels (solid curves) with good and bad behaviour respectively. Good score drops with bad behaviour, and bad score gets better with good behaviour (dotted curves). There is also a score expiring mechanism, which 'decays' good reputation to positive default and bad reputation to negative default, respectively, over time. This is represented using quadratic decay patterns.

The local scores are globally aggregated for collaborative inferences. A global score can initialise a local score if no prior local score is available. The local scores are securely submitted to and stored in a *score aggregation system*. The global score of a consumer is calculated, at query-time, over a window of scores submitted by providers. The impact of each submission is weighed by the rank of the provider. Newer score submission from the same provider overrides the corresponding previous submission. There is time decay, which helps scavenging older scores.

We are investigating an interesting application of the declarative mapping of terms from the service contract as well as the reputation scheme on a front-end load balancer to filter network connections to back-end service endpoints.

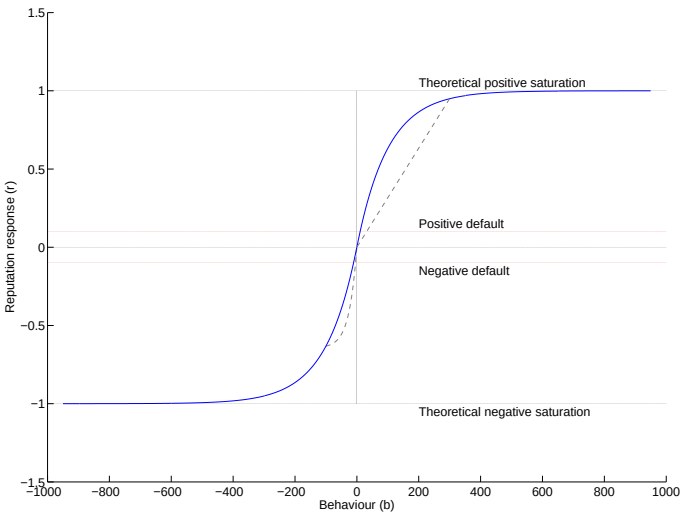

Fig. 1. Graph of local reputation response to behaviour

We are developing a long-lived identity mechanism [4] and also exploring the use of public key infrastructure in the context of our proposed model. We are looking into formal ways of developing provider ranking. The model presented here will be simulated with test data-set from real network logs (e.g., The Internet Traffic Archive[1]). We are looking into an interesting experiment on how to use the consumer score as an incentive mechanism in a peer-to-peer content distribution system. We are investigating the defense against possible attacks (e.g., denial of service, source spoofing) on the model.

References

1. Garriss, S., Kaminsky, M., Freedman, M., Karp, B., Mazières, D., Yu, H.: Re: Reliable Email. Proceedings of the 3rd Symposium of Networked Systems Design and Implementation (NSDI '06) (2006)
2. Golbeck, J., Hendler, J.: Reputation Network Analysis for Email Filtering. Proceedings of Conference on Email and Anti-Spam (CEAS) (2004)
3. Allman, M., Blanton, E., Paxson, V.: An Architecture for Developing Behavioral History. Proc. Workshop on Steps to Reducing Unwanted Traffic on the Internet (2005)
4. Wakeman, I., Chalmers, D., Fry, M.: Reconciling privacy and security in pervasive computing: The case for pseudonymous group membership. Submitted for publication (June 2007)

[1] http://ita.ee.lbl.gov/