School of Computing Science,
University of Newcastle upon Tyne

Choosing reputable Resources in Atomistic Peer-to-Peer Environments

Georgios Pitsilis and Lindsay Marshall

Technical Report Series

CS-TR-816

September 2003

Published by the University of Newcastle upon Tyne,
School of Computing Science, Claremont Tower, Claremont Road,
Newcastle upon Tyne, NE1 7RU, UK.

Choosing Reputable Resources in Atomistic Peer-to-Peer Environments

Georgios Pitsilis[1] , Dr. Lindsay Marshall

University of Newcastle upon Tyne, U.K.,
School of Computing Science
{Georgios.Pitsilis,Lindsay.Marshall}@ncl.ac.uk

Abstract. Peer-to-Peer information sharing environments have gained recognition and popularity during the recent years. In spite of the useful characteristics they provide in the ways that the participants can collaborate, the issue of quality preservation in the shared material has not been addressed yet. The lack of appropriate mechanisms and policies to evaluate the participants has sown fears that the overall popularity of the services will be affected. The nature of atomistic p2p models, where survivability is based on the idea of self-organization into communities could be the basis of a solution to the quality problem build-up by the peers themselves. We consider that the deployment of an assessment scheme as a consultancy service based on a localized view of reputation could help the associated members of the peer-to-peer community in making their choices and thus in the provision of better services.

Keywords: Atomistic Peer-to-Peer, Reputation Management, Word-of-mouth, Electronic Marketplaces, Network Distances.

1 Introduction

The term peer-to-peer has always been used to describe a connectivity architecture which in simple terms means equal communicating with equal. In such a network all the participants offer the same services and obey the same rules. The success of content sharing applications like Gnutella and Free-net has shown that this type of decentralization is likely to have a prosperous future.

2 The problem

The issue of preserving QoS on services of the "static internet" is the only area that has been studied so far [1]. The global form of the contemporary peer-to-peer type of collaboration where every member within the infosphere is freely allowed to share any kind of material raises the issue of how the quality of service can be preserved in

[1] Scholar of Greek State Scholarships Foundation (IKY)

such a model of networking with no central consultancy service. The difference and also the difficulty in p2p systems is found in the dynamic way the peer communities are structured and especially in the *Atomistic* model where the peers cannot exist alone without mutually supporting each other. Unlike the static web, the provision of low quality services has a large impact on performance and can even lead to the collapse of the system [2] when the users receive dissatisfactory services.

The problem comes from the nature of the peer-to-peer model where virtual networks are built upon the collaboration of computers around the world. In this form of networking the nodes/peers are running mostly on ordinary workstations whose connectivity behavior is in fact unstable and unpredictable. This is either due to the use of unreliable transport media or it depends on the working habits of the users. These impact each workstation's ability to serve requests. For example a node that runs on a mobile host is less capable of serving a request due to the higher probability of going off-line during a retrieval operation and thus disrupting the service.

It is common practice on the static web today for similar situations to be addressed by the use of a centralized consultancy services like UDDI [3], or in the simplest way by meta-search engines, which attach as metadata semantic information to search responses. Even though these mechanisms are widely used, they also carry the weaknesses of centralization, (single point of failure, bottlenecks), and hence are unsuitable for P2P environments. The weakness such solution carry is that of showing bias in their judgments towards a particular option. This suggests the use of a decentralized solution.

Reputation based schemes in peer-to-peer networks have also been proposed in the past [4]. The proposed techniques are mainly based on polling mechanisms, which perform reliability checks on candidate participants prior to downloading by sending messages to all neighboring peers. Even though there are no evaluation results for the level of improvement this algorithm provides, the polling mechanism itself is in fact a centralized idea of communicating and obviously impacts the scalability of the system.

3 Our Proposed Solution

Our concept is to provide a solution based on techniques similar to those that are used in everyday life. This is based on the observation of the many similarities between p2p networks and social communities, especially in the ways they grow and exchange information, and the trust relationships build between their entities. In particular, whenever choices have to be made, people seriously consider factors such as *Reputation* and *Trust*. That happens because people tend to trust reputable services and their friends when they ask for opinions rather than other total strangers. Even though reputation theory is quite old as a research issue, it has been applied in computer science for the first time only recently. Reputation in general means "what

we expect about an entity's behavior based on observations and collected information from the environment regarding the past actions of the entity"[5]. We believe that applying a reputation scheme in p2p similar to the one that has been developed within the human communities may benefit system performance by improving the success of queries and at the same time avoiding the problems of centralization.

Our proposed solution is to employ the "word of mouth" techniques to help each peer in the network to build and maintain reputation profiles of how they perceive other peers and to share this information for the benefit of the whole community. "Word of mouth" is a known recommendation mechanism that people use for information diffusion and is basically the way that searching is carried out throughout social relationships. The idea is, people who have lived through an experience before can be a good source of information to others who might want to live the same experience in the future.

Newer organization schemes suggest that peers should get organized into groups based on their interests [6] in the same way that people rely on others who have similar interests. Even though this kind of organization has already been adopted in the newer p2p implementations [7], the information that is exchanged between the nodes (e.g. historical data of past behavior) is not used for the purpose of building reputation profiles.

In order to automate the above mechanisms of building social relationships, we consider the virtual community as a directed graph with weighted edges that correspond to the reputation ratings that one node gives to another. The graph grows as the experiences between the nodes increase and more reputation ratings take place. Under this scheme, each node, before choosing a peer, could gather via "word of mouth" enquiries as much information as possible, to build the reputation profile of the target peer. A typical reputation enquiry could be initiated via an ordinary broadcast, hopping from node to node recursively until it reaches the nodes close to the examined one. Any direct links with the target could denote the existence of personal experiences and thus their ability to providing personalized reputation ratings. The neighbors in the scheme act all together as a formed group and actually carry out this "survey" by providing opinions. It is worth pointing out that not all neighbors' opinions are given the same weight during the assessment, but the weight value is given in proportion to the reputation measure that corresponds to the "path" where the information comes from. The notion of intermediate reputation ratings offers a quite strong characteristic, as the actual reputation value can be determined considering all the intermediate nodes that carry the responses back to the peer where the query was originated. Likewise, in the real world example, the opinions are also freely formed. To develop a human like judgment for the model, we could re-rate the neighbor's reputation values, after a choice has been made, by giving new weights depending on how close the suggested reputation had been to the actual one.

In contrast to the traditional centralized reputation systems that require the

existence of a central trusted entity as a deposit for the behavioral data (e.g. UDDI), our algorithm requires no central entities at all to store the reputations. Instead, that information is gathered dynamically whenever reputation queries are fired inside the system graph and the trust relations with other peers are built as the originating peer evaluates the reputation ratings.

The model we described for building reputation profiles employs operations analogous to the Histos [8] system, which was designed especially as a collaborative reputation mechanism for electronic marketplaces. That system was inspired by *the Friend of a Friend Finder* Scheme [9] where it is thought that any persons within a community can be known to each other through their relationships with their friends.. Respectively, any relationships established between entities could be based on the commonalities that can be found on their sets of interests.

We believe that this algorithm, if applied to P2P systems, will show powerful characteristics, such as the ability to generate reputation dependent on the point of view, which is a fundamental property of P2P. Nevertheless, P2P networks use for their connectivity a virtual network that links the nodes together, and not considering this along with the reputation norm would be a serious omission. Given that a P2P network is substantially an application layer network built upon the Internet infrastructure, in our case, is desirable to know the physical network distances between the peers because the actual service provision is done at this level. Thus, a decision based on a hop distance is not always correct and it can lead to wrong estimates of a peer's "reputability". This is because the virtual network topology, as various studies about the Gnutella network infrastructure [10] have shown, mismatches the Internet topology where Gnutella runs. Changing the topology as a remedy is not always applicable and is prohibitively expensive. Instead, we could focus on the ways that gathered information is assessed and to consider the topology information to some extend in rating the reputability of a peer. That way, in our proposed scheme, the topology information is upgraded to an influential factor offering also personalized characteristics in the shaping of reputation.

4 Conclusion

In conclusion, we believe that the way that Reputation is built in human social communities provides a good analogy to studying the problems of quality of services in Atomistic peer-to-peer networks. The numerous benefits of distributed computing such as congestion avoidance and the non-existence of a single point of failure, that are met in this model, combined with the characteristics that are found on the human trading societies, are promising enough to build a useful platform. Moreover, considering that the underlying network infrastructure offers a strong dependency on the user's point of view, that is otherwise missing, it can provide enhanced characteristics and get the model closer to a real situation.

References

1. Raju Pandey - J.Fritz Barnes, "Supporting QoS in HTTP servers", Proceedings of the Seventeenth Annual SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC). (Puerto Vallarta, Mexico, June 1998).
2. H.Chen, A.M.Schroeder, "A Modified Depensation Model of Peer to Peer Networks: Systematic Catastrophes and other potential Weaknesses", AMATH 383 – June 2002
3. Universal Description Discovery and Integration Technical White Paper. September 2000
4. E.Damiani – S.DeCapitani –S.Paraboschi, P.Samarati,"A Reputation-Based approach for finding Reliabe Resources in Peer-to-Peer Networks", CCS'02 Nov 18-22, 2002 Washington DC, USA
5. A.Abdul-Rahman – S.Heiles, "Supporting trust in Virtual Communities", proceedings of Internatinal conference On System Sciences, Jan 4-7 2000, Hawaii
6. K.Sripanidkulchai, B.Maggs, H.Zhang, "Efficient Content Location Using Interest Based Locality in Peer-to-Peer Systems", Infocom 2003.
7. M.Prinkey, "An efficientscheme for query Processing on Peer-to-Peer Networks", Aeolus Research Inc.
8. G.Zacharia - A.Moukas - P.Maes, "Collaborative Reputation Mechanisms in Electronic Marketplaces", Proceedings 32th Hawaii International conference on system Sciences – 1999
9. N.Minar- A.Moukas, "Friend of a friend finder", Software agents Group, MIT Media Lab, working paper 1988.
10. Matei Repeanu, "Peer to Peer Architecture case study: Gnutella Network", University of Chicago, Dept. of Computer Science, Tech Report TR-2001-26