Score Manager Discovery in EigenTrust Using Virtual Magnetic Fields

Henrique Galperin, Luiz Augusto de Paula Lima Jr. and Alcides Calsavara
Pontifícia Universidade Católica do Paraná (PUCPR)
R. Imaculada Conceição 1155 - Curitiba - Paraná - Brazil
Email: {henrique.galperin,laplima,alcides}@ppgia.pucpr.br

***Abstract*—EigenTrust is a well-known distributed reputation system that uses random nodes (called Score Managers) to compute the reputation of other nodes in the network. In the original proposal, the selection of the Score Managers is made by successive hashes in a DHT. This paper proposes the usage of Virtual Magnetic Fields in replacement of DHTs in the selection of Score Managers in the EigenTrust reputation system. Virtual Magnetic Fields are self-organizing message forwarding mechanisms that are capable of delivering messages to specific nodes according to some non-functional aspects concerning the application semantics. A comparative analysis showed that the EigenTrust efficiency was improved in the proposed solution by removing the cost of the discovery of a Score Manager, and by introducing a low cost method for dissemination of information as nodes join and leave the network.**

***Keywords*-EigenTrust; reputation systems; P2P; virtual magnetic fields.**

I. INTRODUCTION

Peer accountability has long been a problem to peer-to-peer networks [1]. In order to address the issue, several distributed reputation systems have been proposed in the literature. Among them, we find EigenTrust [1], which is a reputation system that assigns to each peer a unique global trust value, based on the peer's history. It uses a distributed and secure method to compute global trust values for each peer. Those values can be used by peers to filter the interaction with other peers.

In EigenTrust, the trust value of each peer is computed by a collection of other nodes called Score Managers. The Score Managers of each peer node are randomly located by successively applying a hash function of a unique ID of the peer, such as its IP address and TCP port, resulting in a point in a DHT (Distributed Hash Table) hash space.

This paper proposes an improvement of EigenTrust replacing DHTs by Virtual Magnetic Fields for the selection of Score Managers. A Virtual Magnetic Field is a distributed self-organizing message forwarding mechanism based that allows routing of message to the most "attractive" nodes based on an attraction force function created according to the application semantics.

This paper is organized as follows. Section II details the EigenTrust reputation system and presents the Virtual Magnetic Fields distributed message routing paradigm. After discussing the motivations for this work, Section III introduces the topology establishment and maintenance algorithm and explains how to compute the attraction strength of each node. Then, a node grouping methodology is presented whose goal is to optimize virtual magnetic planes. The proposed solution then is compared with the current solution in Section IV. Finally, future works are discussed and conclusions are drawn in Section V.

II. RELATED WORK

In this section, the two main related works are presented, namely, the EigenTrust distributed reputation system and the the Virtual Magnetic Networks.

A. EigenTrust

EigenTrust [1] is a distributed reputation system for peer-to-peer networks. It is based on the idea that each peer in the network must have a global reputation value that reflects the experience that all nodes in the network had with it. The reputation value is first computed locally by normalizing the difference between the number of positive transactions and the number of negative transactions with a peer in order to obtain a result between 0 and 1. The global reputation of each node is computed based on the notion of transitive trust, by weighting the local trust value assigned to a node with the global reputations of it provided by other peers.

For a secure version of the reputation system, the paper proposes to use DHTs, like CAN (Content Addressable Networks) [2] and Chord [3], and to apply a hash function to the unique Id of some node x (e.g., x's IP address and TCP port) in order to map it into a point in a DHT hash space. The peer which currently covers this point will be considered the Score Manager of x. Successive hashes of the node's Id are used to determine the others Score Managers for that node. The Score Managers are responsible for consolidating the reputation of a given node.

B. Virtual Magnetic Networks

Message routing based on specific application requirements has emerged as an interesting research field due to the fast growing domain of distributed applications, especially where mobile platforms are employed. A novel message routing mechanism was introduced in [4] and [5], where the concept of *magnetic fields* is borrowed from physics to define node relationships which permit a node to attract messages, like a magnet attracts iron. Such mechanism suits applications that

Copyright (c) IARIA, 2011. ISBN: 978-1-61208-161-8

require messages to be delivered according to some particular non-functional requirement, that is, it can be customized according to some application needs. As an example, an application where messages carry tasks to be performed may require that every single message should be delivered to the current network node where the corresponding task can be performed in the shortest time. Naturally, that depends on node processing capacity, which changes dynamically and can be hard to manage. By employing the routing mechanism based on magnetic fields customized to attract messages to the fastest node, such application is released from the burden of managing node processing capacity information.

The magnetic-field-based routing mechanism employs an *overlay* network (represented as a directed graph), which hides the physical network, in order to model each particular application's non-functional requirement accordingly. Moreover, the mechanism permits logical node mobility, thus changing network topology, that is, an overlay network is dynamic. A *virtual magnetic network* is defined as an overlay network where each node contains a virtual magnet that attracts messages. Each virtual magnet has a *force* that represents some property associated to the target application's non-functional requirement. Thus, a node – through its virtual magnet – magnetically influences its neighbors, such that a message can be attracted from a neighbor. Magnetic influence relationships are transitive, such that a message can be attracted from an indirect neighbor, as well. The main goal is to deliver each message to the strongest node – named *pivot* – according to magnetic influence relationships, as defined by the corresponding overlay network graph, independently of the node where a message is firstly created. The mechanism assumes that each node determines its own force and, as it changes, publishes its new force to neighbors.

III. Score Manager Discovery Using Virtual Magnetic Fields

Although EigenTrust, in its original specification, uses a DHT for the selection of Score Managers, we suggest that it can be replaced by a Virtual Magnetic Network decreasing the network load in Score Managers lookup operations.

The main idea is to use one virtual magnetic plan per Score Manager. The pivot node in a plan will be the selected Score Manager. The magnetic forces of every peer are random and auditable by any other peer in the plane, to make sure that the chosen Score Manager is really random and has not been manipulated.

Since the reputation of each node can be assessed by s Score Managers (typically a global constant value) and since each Score Manager requires a particular attraction plane, we end up with $s \times N$ planes (N is the number of nodes in the network). In order to reduce this number, it is possible to group peers, allowing them to share the same set of Score Managers.

A. Motivation

The main motivation for the replacement of DHTs for Virtual Magnetic Networks is that Virtual Magnetic Networks have a proactive way of handling routing, while DHTs have a reactive way. This means that when a peer needs to know who is the Score Manager of another peer, when using a DHT, the peer will need to make a lookup in the network, contacting some nodes in the way to have the answer. This is not true when using Virtual Magnetic Networks, since the information is, in a proactive way, already known by the peer.

Even having a proactive algorithm, as proofed in the next sections, this does not cause a big overload in the network on the joining or leaving of a peer form the network, since there will only be a big number of messages exchanged when the Score Manager changes.

That fact makes the EigenTrust algorithm much more efficient in its most common operations, like the lookup of a Score Manager and the joining or leaving of a non Score Manager node, and just a few slower in the lees used operations, like the joining or leaving of a Score Manager.

B. Establishing and Maintaining the Topology of the Virtual Magnetic Networks

In order to replace DTHs by Virtual Magnetic Networks in EigenTrust, one must define how the peers inside the virtual magnetic planes are organized. Two main plane topology categories can be identified: the static topology, where there is no joining or leaving of nodes, and the dynamic topology, with nodes joining and leaving the plane at any moment.

In a static topology scenario, a good option is to construct the planes as Small Worlds [6] with a reduced number of edges, but at the same time, without increasing the average number of hops between nodes. Other techniques can be used as well, including the reproduction of the actual underlying physical topology.

When considering a dynamic topology, there is always the risk of disconnecting the virtual magnetic plane, when nodes leave the network. However, this disconnection does not prevent EigenTrust from working, but may compromise only its performance since each separate plane would have its own Score Manager. Nevertheless, all the Score Managers would still compute the same global trust value. If this situation is not avoided or minimized, the segmentation of planes tends to grow over time, resulting in an undesirable number of Score Managers. So some technique is needed to avoid plane segmentation in a dynamic network topology. Assuming that each node is capable of detecting the disconnection of a neighbor, a polling mechanism may be set up in order to minimize the probability of segmentation. A joining node x just needs to get connected to a known bootstrap peer. It is desirable, on the other hand, that x should get connected to at least some other peer chosen randomly, in order to reduce the chance of splitting the plane in case of later disconnections. If x is the new pivot, then this information is disseminated using the traditional Virtual Magnetic Networks algorithm [5].

There are two categories of nodes that can leave the network: pivots and "regular nodes" (i.e., non pivots). If a regular node leaves the magnetic plane, then the network will be affected only if that node is part of the route from another

Copyright (c) IARIA, 2011. ISBN: 978-1-61208-161-8

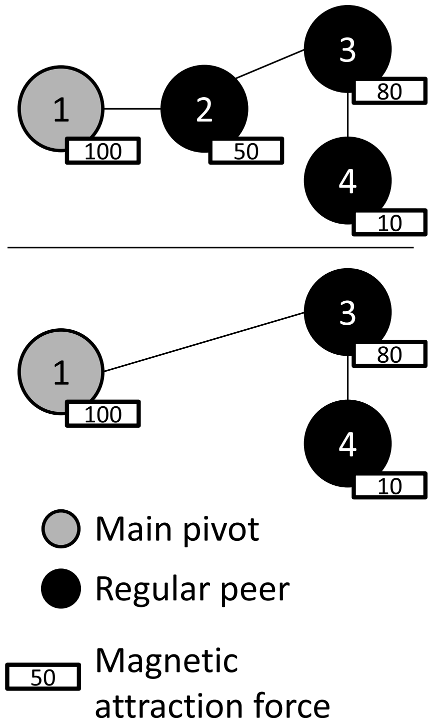

Figure 1. Leaving of a regular node that belongs to the route to the pivot.

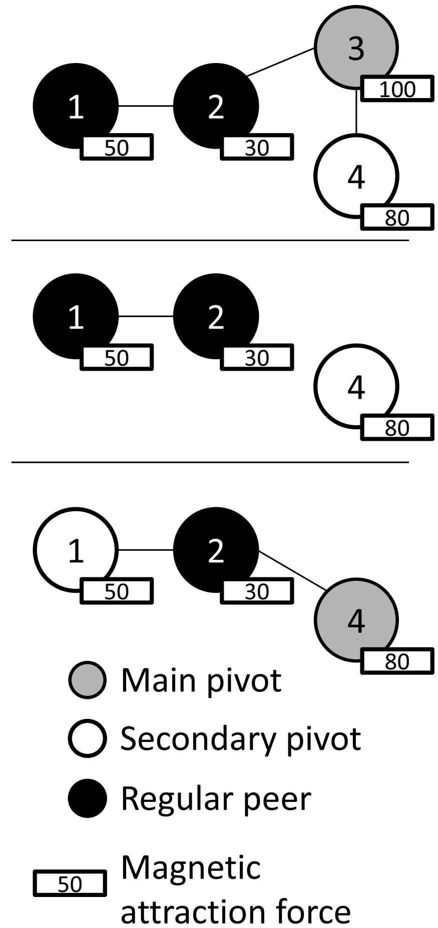

Figure 2. Leaving of the pivot.

regular node to the actual pivot. When a regular peer x detects that one of its neighbors x_n has left, that peer should check whether the missing node x_n is in its route to the pivot. If this is not the case, then nothing needs to be done. If, however, x_n is in the path to the pivot, then x will get connected directly to the pivot in order to keep the network cohesive and to restore its connection to the pivot. This is always possible because every peer knows the identity of the pivot. Since in this case the pivot has not changed, all the preexistent routes will still be valid, and no attraction force change needs to be propagated. This process is depicted in Figure 1, where node 2 leaves the plane and forces node 3 to establish a new connection with the pivot.

On the other hand, if the leaving node is the pivot, then a secondary pivot (i.e., the peer with the second largest attraction force in the plane) becomes the new pivot and its attraction strength is propagated to all nodes in the plane. If it leaves the network before the pivot itself, then the same procedure used when a regular node leaves is executed, but now a new secondary pivot has to be found.

Therefore, when a pivot leaves the network, the secondary pivot becomes the pivot, and all the neighbors of the leaving pivot check if they still have an active route to the new pivot. If they do not, they get connected directly to the new pivot. This procedure guarantees the cohesion of the plane topology. Since the pivots in the plane have changed, all attraction forces need to be propagated, starting by the nodes that lost direct connection to the previous pivot. Naturally, the virtual magnetic plane will elect a new secondary pivot and will be updated with the new pivot, using the basic force propagation algorithms of Virtual Magnetic Networks [5].

Figure 2 depicts a scenario where the pivot (node 3) leaves the network. After that, node 2 gets connected to the secondary pivot (node 4), which will become the new pivot and then a new secondary pivot is elected (node 1, in the example).

C. Computation of Attraction Forces

Once the magnetic network topology has been established, the next step consists in defining a method to assign attraction forces to each node. Since this attraction force will be used to select the score manager for a given node (or node group), it must be unique and random within an attraction plane and different across the multiple planes.

In order to guarantee anonymity and randomness required by EigenTrust, the attraction force $F(x)$ of a given node x is defined by Equation 1.

$$F(x) = H(H(I(t)) + I(x) + k) \qquad (1)$$

where,

H is a well-known and reliable hash function, such as SHA1 [7];

$I(n)$ is the identifier of node n;

t is the node whose reputation is evaluated by the score manager;

k is a natural number used to distinguish score managers belonging to different attraction planes. For instance, if there are three score managers for each node (or group of nodes) evaluated, then k may range from 0 to 2.

Equation 1 is random due to the hash function H, guarantees anonymity by using $H(t)$ instead of $I(t)$ directly, and minimizes the probability of existing equal forces for different nodes within the same plane, since $I(n)$ is distinct for each n. Moreover, the probability of having the same score manager in different planes is minimal since k takes different sequential values.

Notice as well that all parameters of Equation 1 are known to all nodes in the magnetic plane. As a consequence, it is virtually impossible for a malicious node to fake its attraction force and to corrupt attraction force propagation data.

D. Peer Grouping

From Equation 1, one can observe that a magnetic plane can be identified by constant factors that are the same for all nodes in the plane, namely, the pair $[H(I(t)), k]$. Therefore, if we multiply the number of existing $H(I(t))$ by the number of planes (i.e., score managers) per node, we obtain the total amount of planes in the network.

Note that if there is a different $H(I(t))$ for each t (t, is a node in the network), the total number of planes can be very high, and this brings overhead to the network. In order to reduce the number of planes, multiple nodes can be grouped so that they use the same planes and the same score managers. This is done by performing the integer division of $H(I(t))$ by a constant g, for each node t. Nodes that exhibit the same result will belong to the same group, and $H(I(t))/g$ will replace $H(I(t))$ in Equation 1. For example, if there are three nodes x, y and z in a network, and $H(I(x)) = 20$, $H(I(y)) = 24$ and $H(I(z)) = 35$, then, if $g = 10$, nodes x and y will belong to the same group (since $H(I(x))/g = H(I(y))/g = 2$) and, therefore, they will share the same score managers. This strategy reduces the risk of manipulation by creating random groups.

Notice that it is important to choose the constant g according with the expected number of nodes in the network and the magnitude of H. For instance, if there are just a few nodes in the network and H produces very large numbers, it is desirable to use a very large value for g, so that to increase the probability of two nodes belong to the same group.

E. Virtual Attraction Forces Propagation

Considering that the attraction force of every peer in the magnetic pane can be calculated by any other peer, it is possible to simplify the force propagation algorithm. It is possible to propagate only the peer id, removing the attraction force from the propagation tuple, and compute the peer force locally.

IV. Comparison with the Current Solution

A comparison between magnetic field networks and DHTs as alternative approaches to implement a system to determine score managers based on EigenTrust can be made by assessing network and node resources usage in both cases. The following events must be taken into account to make a fair comparison: node join into the network, node leave from the network, and selection of score manager for an arbitrary node.

If only score manager selection is considered, it is possible to notice that the approach based on magnetic field networks presents a clear advantage, since all nodes know the pivot all the time and, consequently, no messages are needed to select the score manager, while in the approach based on DHTs, a navigation through the network is needed in order to find out which node occupies the position corresponding to the score manager, thus causing several message exchanges.

In the event of node entry and exit, there are only two cases which may cause some performance impact for the approach based on virtual magnetic networks. Firstly, if a node that joins or leaves the network is the score manager, there will be a some performance degradation, since an update regarding pivot information will be triggered. In all other cases, the consequences are not relevant. On the other hand, if for instance CAN is employed, the impact is almost always low, since in most cases, only neighbor nodes need to be notified in order to either split the existing area (when a node enters the network) or occupy a newly empty area (when a node leaves the network) if no take-over is needed. Hence, the approach based on DHTs performs better than the approach based on virtual magnetic fields in the cases where a node that enters or leaves the network is the score manager. However, the larger network, the lowest are the chances for that to happen.

Therefore, assuming that score manager join and leave happen in a much lower rate than score manager selection, the approach based on magnetic field networks will have a better overall performance than the approach based on DHTs. The costs of node entry and exit in virtual magnetic networks and in CAN are analyzed in the following sections.

A. Message Cost Analysis in Virtual Magnetic Networks

Consider the following variables:

N the number of peers in the plane;

E the number of edges (connection between two peers) in the plane, assuming that all edges are bidirectional;

Ei the number of edges created by a peer that is joining the network;

Ce the cost of creation of a new bidirectional edge;

Ea the average number of edges per peer.

The probability (P) of a peer that is joining or leaving the magnetic plane is a primary or a secondary pivot can be defined by Equation 2.

$$P = \frac{2}{N} \qquad (2)$$

If a pivot changes, the number of messages needed to select a new pivot is given by Equation 3 (cost of propagation).

$$Cp = 2 \times E - N + 1 \tag{3}$$

Based on Equation 3, it is possible to calculate the peer joining average cost (Cja) in a plane (Equation 4).

$$Cja = Ei \times Ce + P \times Cp \tag{4}$$

Let Pr be the probability of an edge connected to a leaving node is part of a connection route to the pivot node of the node in the other side of the edge, then the average cost of a peer that leaves the network (Cla) is given by Equation 5.

$$Cla = Ea \times Pr \times Ce + P \times Cp \tag{5}$$

Considering that the cost of creation of a new bidirectional edge (Ce) is generally low, the peer joining or leaving the network, in terms of number of messages exchanged ($Cjla$) can be approximated by Equation 6.

$$Cjla = \frac{4 \times E + 2}{N} - 2 \tag{6}$$

Since the total number of edges (E) is function of the average number of edges per peer (Ea) and the number of peers in the plane (N) as shows the Equation 7, then $Cjla$ can be reduced to Equation 8.

$$E = \frac{Ea \times N}{2} \tag{7}$$

$$Cjla = \frac{2}{N} + 2 \times Ea - 2 \tag{8}$$

Notice that as N increases the peer joining and leaving average cost ($Cjla$) tends to be influenced only by the average number of edges per peer (Ea), which is independent of the network size. This result indicates that the solution proposed scales. Moreover, Ea also tends to be small, since it depends mainly on the number of edges created by a peer that is joining the network (Ei). Obviously this parameter that can be set as low as 2 and generally there is no good reason to set it to a larger value.

Notice as well that even considering that the cost of creation of a new edge on the graph (Ce) greater then zero, the final result would also be independent of the number of peers (N), which confirms the conclusion we reached.

B. Lookup Rates and Session Lengths

Regarding the lookup of score managers, the use of Virtual Magnetic Networks is clearly advantageous over traditional DHTs, since there is no need of communication, while DHTs require $O(logN)$ at best (CHORD). The information about score managers is already known by each node due to the proactive nature of the magnetic force propagation algorithm.

In the case of Virtual Magnetic Networks, all costs are transferred to the moment nodes join or leave the network. Although DHTs may show a constant cost in these situations (generally speaking, leaving nodes require updates of routing tables), notice that in $(1 - 2/N) \times 100$ percent of the cases Virtual Magnetic Network will require just a few reconnection operations. Only when the node joining or leaving the network is the pivot (primary or secondary) then propagation will be required, and the message cost will be higher. In fact, if we consider the use of CAN to implement the DHT abstraction (actually, this is the P2P (Peer-to-Peer) network suggested by the EigenTrust original specification [1]), the leaving process can be even more expensive, since it is possible that none of the neighbors of the leaving area can occupy the empty space, forcing a "take-over situation", that has no time upper bounds.

Since typical session lengths (a session is defined as join-participate-leave cycle) in P2P structured networks can be measured in hours as shown in [8] and since the probability of a pivot joins or leaves the network is small $(2/N)$, we can safely claim that the lookups outnumber by far the need of propagation due to pivot changes. This fact makes the use of Virtual Magnetic Fields more advantageous than DHTs in this context.

V. Conclusion and Future Work

We have presented an alternative to DHTs for the selection and lookup of Score Managers in the EigenTrust reputation system based on Virtual Magnetic Networks. The proposed solution provides a proactive solution for this problem, without adding significant costs as peers join and leave the network. Specifically, the method removes the need of peer lookup in order to identify the Score Managers of a particular node. As shown in Section IV, the use of Virtual Magnetic Fields in this context brings a real gain in the average network operation, providing a real gain in the average network operation.

This work is part of an ongoing research whose goal is to define a complete reputation mechanism for Virtual Magnetic Fields in an open network. Our future work therefore includes the design and implementation (possibly through simulation) of a reputation system that could weight the attraction strength informed by each node based on the reputation of that node.

References

[1] S. Kamvar, M. Schlosser, and H. Garcia-Molina, "The eigentrust algorithm for reputation management in p2p networks," in *Proceedings of the 12th international conference on World Wide Web*. ACM, 2003, pp. 640–651.

[2] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, "A scalable content-addressable network," in *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2001, pp. 161–172.

[3] I. Stoica, R. Morris, D. Karger, M. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2001, pp. 149–160.

[4] L. Lima and A. Calsavara, "Autonomic Application-Level Message Delivery Using Virtual Magnetic Fields," *Journal of Network and Systems Management*, vol. 18, no. 1, pp. 97–116, 2010.

[5] A. Calsavara and L. Lima Jr, "Routing Based on Message Attraction," in *2010 IEEE 24th International Conference on Advanced Information Networking and Applications Workshops*. IEEE, 2010, pp. 189–194.

[6] D. Watts and S. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.

[7] P. FIPS, "180-2: Secure Hash Standard," *US Department of Commerce, Technology Administration, National Institute of Standards and Technology*, 2002.

[8] D. Stutzbach and R. Rejaie, "Understanding churn in peer-to-peer networks," in *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. ACM, 2006, pp. 189–202.

Copyright (c) IARIA, 2011. ISBN: 978-1-61208-161-8