CS 856 Paper Reviews
W Anthony Young - 20161423
November 18, 2004

Paper Title: The EigenTrust Algorithm
Section I. Overview

A. Reader Interest

1. Which category describes this paper?

( ) Practice / Application / Case Study / Experience Report
(X) Research / Technology
( ) Survey / Tutorial / How-To

2. How relevant is this paper to the readers in the area? Please explain your rating under III. Detailed Comments.

(X) Very Relevant
( ) Relevant
( ) Interesting - but not very relevant
( ) Irrelevant

B. Content

1. Please explain how this paper advances this field of research and / or contributes something new to the literature. Please explain your answer under III. Detailed Comments.

2. Is the paper technically sound? Please explain your answer under III. Detailed Comments.

( ) Yes
( ) Appears to be - but didn't check completely
(X) Partially
( ) No

C. Presentation

1. Are the title, abstract, and keywords appropriate? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) No

2. Does the paper contain sufficient and appropriate references? Please explain your answer under III. Detailed Comments.

(X) References are sufficient and appropriate
( ) Important references are missing; more references are needed
( ) Number of references are excessive

3. Does the introduction state the objectives of the paper in terms that encourage the reader to read on? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) Could be improved
( ) No

4. How would you rate the organization of the paper? Is it focused? Is the length appropriate for the topic? Please explain your answer under III. Detailed Comments.

(X) Satisfactory
( ) Could be improved
( ) Poor

5. Please rate and comment on the readability of this paper.
Please explain your answer under III. Detailed Comments.

( ) Easy to read
(X) Readable - but requires some effort to understand
( ) Difficult to read and understand
( ) Unreadable

Section II. Summary and evaluation

A. Summary (provide here a summary of the paper)

This paper presents a method for managing global trust of peers in a P2P network. They assign global trust values that are computed securely. The authors begin with a description of the problem and motivation to find a solution.

The authors present the five main design considerations that must be taken into account when developing a trust and reputation system, such as anonymity of peers and minimal overhead of administration.

The authors turn their attention to discussion of how trust values can be managed by a central server, followed by a distributed secured trust management system. In this manner, it is supposedly not possible for malicious peers to present poor trust scores for other peers, and good trust scores for itself. A discussion of how trust values can be used (such as isolating malicious peers and providing incentives to freeloading peers to participate) is presented.

The authors proceed by giving a performance evaluation. The EigenTrust method converges to an accurate trust value after only a few iterations of the trust calculation. Probabilistic trust downloading makes sure that load is distributed among several good peers, and allows new peers to build up reputation slowly. Different threat models are presented and the algorithms are discussed against those threat models.

The authors conclude with a discussion of the related work and some summary comments.

B. Evaluation

Please rate the paper. Please explain your answer under III Detailed Comments.

( ) Award Quality
(X) Excellent
( ) Good
( ) Fair
( ) Poor

Section III. Detailed Comments

Overall, you have presented a novel concept for managing reputation in a system. Your presentation is well written, but contains some grammatical errors and missing words. As well, your presented mathematical models are vaguely explained and take considerable effort from readers to understand.

Your introduction is good, as it outlines the nature of the problem and, briefly, your solution. The abstract, title, and keywords are appropriate to the content, but you should have included “trust management” or “trust” in the keywords to more properly identify the content of the paper.

It appears that this work will further the field of trust and reputation. This distributed algorithm is quite novel. It appears to be one of the first to be presented, but requires some additional questions to be answered:

-Is it possible that peers will submit only negative reputation values even if the peer deserves a positive reputation value? How do you overcome this problem when one malicious peer can conceivably create many different identities each submitting only negative trust values for each user requested? I am not convinced that the captcha approach is enough to stop a Sybil Attack as users only need time to read and generate their new identities?
-Is it possible that malicious peers can collude to provide each other with positive trust values even though they do not deserve them? I am not convinced that your method of providing some probability to the pre-trusted peers is enough to stop this from happening.
-No discussion of scalability of this system is given. How well does your algorithm scale? Many peer networks have hundreds of thousands of nodes, and your simulations use a maximum network size of just over 100!

Clearly, you have attempted to cover many of the details of your algorithm. However, the above should be answered as well.

Paper Title: Automated Trust Negotiation
Section I. Overview

A. Reader Interest

1. Which category describes this paper?

( ) Practice / Application / Case Study / Experience Report
(X) Research / Technology
( ) Survey / Tutorial / How-To

2. How relevant is this paper to the readers in the area? Please explain your rating under III. Detailed Comments.

(X) Very Relevant
( ) Relevant
( ) Interesting - but not very relevant
( ) Irrelevant

B. Content

1. Please explain how this paper advances this field of research and / or contributes something new to the literature. Please explain your answer under III. Detailed Comments.

2. Is the paper technically sound? Please explain your answer under III. Detailed Comments.

(X) Yes
( ) Appears to be - but didn't check completely
( ) Partially
( ) No

C. Presentation

1. Are the title, abstract, and keywords appropriate? Please explain your answer under III. Detailed Comments.

( ) Yes
(X) No

2. Does the paper contain sufficient and appropriate references? Please explain your answer under III. Detailed Comments.

(X) References are sufficient and appropriate
( ) Important references are missing; more references are needed
( ) Number of references are excessive

3. Does the introduction state the objectives of the paper in terms that encourage the reader to read on? Please explain your answer under III. Detailed Comments.

( ) Yes
(X) Could be improved
( ) No

4. How would you rate the organization of the paper? Is it focused? Is the length appropriate for the topic? Please explain your answer under III. Detailed Comments.

(X) Satisfactory
( ) Could be improved
( ) Poor

5. Please rate and comment on the readability of this paper.
Please explain your answer under III. Detailed Comments.

( ) Easy to read
( ) Readable - but requires some effort to understand
(X) Difficult to read and understand
( ) Unreadable

Section II. Summary and evaluation

A. Summary (provide here a summary of the paper)

This paper presents a method of automatically negotiating trust based on an exchange of credentials/certificates. The authors begin with a discussion of the problem as well as a brief outline of their solution.

The authors begin with a discussion of some methods of trust negotiation and introduce their method: finding a sequence of credentials

that can be revealed in order to convince each party that the other is who they say they are. Some related work is presented and the notion of Automated Trust Negotiation (ATN) is formalized using mathematical notation.

The authors next discuss their augmentation of two trust negotiation models using locally trusted third parties (LTTP's) to break a dependent cycle in credentials required by parties. The authors continue by discussing how LTTP tables are maintained, as well as how trust negotiation can benefit from reputation systems.

The authors discuss how their algorithm helps improve the number of successful negotiations of credentials. With an LTTP table of unlimited size, the number of successful negotiations is highly increased. In fact, with a table limited to 10 entries, successful negotiations increases substantially. The authors compare different table negotiation methods.

Some discussion of related work and summary remarks is presented in conclusion.

B. Evaluation

Please rate the paper. Please explain your answer under III Detailed Comments.

( ) Award Quality
( ) Excellent
(X) Good
( ) Fair
( ) Poor

Section III. Detailed Comments

You have presented an interesting concept of using third parties to help mediate trust negotiation. Overall, your writing style is good. Some missing words and grammatical errors make the paper hard to follow in places. Your abstract and introduction are a well written outline of the problem, but do not discuss your proposed method very well. Also missing is a set of keywords to classify your paper.

While I believe that your notions of using a third party intermediary are novel and will advance the field, I am confused with your discussion of the solution. Many of your discussions did not appear to lead anywhere, and were wordy and hard to understand. Much greater care is needed next time

you publish to ensure that your paper is readable by those of varying backgrounds.

I believe your performance evaluation is decent, but could have been improved in a few ways:

-Your work is missing a discussion of how much overhead is added using LTTP's. Does this increase the negotiation time? How long does negotiation take? How much overhead is required to maintain the tables?
-Your work is missing a discussion of the scalability of your algorithm. How well does it scale? You tested your method with larger numbers of nodes, but did not present any of these findings. How does the system behave? Even a brief discussion would have been better than no discussion at all.

I believe the above deficiencies need to be fixed before this paper is to be published, but I do believe you have a novel concept.

Paper Title: Reputation-Based Approach for Choosing Reliable Resources

Section I. Overview

A. Reader Interest

1. Which category describes this paper?

( ) Practice / Application / Case Study / Experience Report
(X) Research / Technology
( ) Survey / Tutorial / How-To

2. How relevant is this paper to the readers in the area? Please explain your rating under III. Detailed Comments.

(X) Very Relevant
( ) Relevant
( ) Interesting - but not very relevant
( ) Irrelevant

B. Content

1. Please explain how this paper advances this field of research and / or contributes something new to the literature. Please explain your answer under III. Detailed Comments.

2. Is the paper technically sound? Please explain your answer under III. Detailed Comments.

( ) Yes
( ) Appears to be - but didn't check completely
(X) Partially
( ) No

C. Presentation

1. Are the title, abstract, and keywords appropriate? Please explain your answer under III. Detailed Comments.

(X) Yes

( ) No

2. Does the paper contain sufficient and appropriate references? Please explain your answer under III. Detailed Comments.

( ) References are sufficient and appropriate
(X) Important references are missing; more references are needed
( ) Number of references are excessive

3. Does the introduction state the objectives of the paper in terms that encourage the reader to read on? Please explain your answer under III. Detailed Comments.

( ) Yes
(X) Could be improved
( ) No

4. How would you rate the organization of the paper? Is it focused? Is the length appropriate for the topic? Please explain your answer under III. Detailed Comments.

(X) Satisfactory
( ) Could be improved
( ) Poor

5. Please rate and comment on the readability of this paper.
Please explain your answer under III. Detailed Comments.

(X) Easy to read
( ) Readable - but requires some effort to understand
( ) Difficult to read and understand
( ) Unreadable

Section II. Summary and evaluation

A. Summary (provide here a summary of the paper)

This paper presents the XRep protocol for automated reputation management. The authors begin with a brief discussion of peer-to-peer systems as well as the architecture of the Gnutella network that was used as the testbed for their implementation.

After the introductory material, the authors present their protocol for reputation management (i.e. XRep). The XRep adds an additional step to the download process where reputations of peers is polled and evaluated before choosing which resource to download.

The discussion then turns to some of the attacks that can be performed on this kind of system (i.e. shilling and man-in-the-middle attacks). The authors discuss how XRep protects against each of these attacks. A presentation of how the combination of servent and resource digests helps in managing reputation is also presented.

The authors then discuss the distribution of resources and users in the Gnutella network. The authors conclude with a discussion of some related work as well as summary remarks.

B. Evaluation

Please rate the paper. Please explain your answer under III Detailed Comments.

( ) Award Quality
( ) Excellent
(X) Good
( ) Fair
( ) Poor

Section III. Detailed Comments

You have presented a good proposal for a distributed reputation based system. Your article is very well written. No spelling and very few grammar mistakes exist. Also, your writing style is appropriate for users of varying technical levels. Your contribution is a new protocol that will be sure to require study by leading researchers. Your article is well organized into self-contained and relevant sections of information.

Specifically, I felt that your introduction did not provide a good enough overview of your protocol. Perhaps more information about how it works should have been included to give the reader a flavour of what is about to be presented. Your presentation is lacking answers to some very important questions (see below). Also, I believe you are missing some important Gnutella distribution information from other papers. Your claims that Gnutella is Zipf appear to contradict other's findings.

Some unanswered questions that I still have are:

-I am not convinced that requiring real IP's is sufficient to stop peers from colluding. Can you explain how this is so? Also, if one knows the reputation policy, is it not possible for them to manufacture the required trueVote messages?
-Your algorithm appears to effectively double the number of messages required to find and download a file. Does this not further degrade the performance of a network already troubled with too many messages?
-Why is it unlikely that D will be chosen for download in the man-in-the-middle attack?
-It appears possible for peers with negative reputation to simply change their ID's and start over. Doesn't this defeat the whole purpose of such a system?
-What is the relevance of discussing the distribution of resources in the Gnutella network, and why do your results contradict those of other papers in claiming that Zipf's law does not hold in these systems?

In total, I believe your paper is good, but that it requires the answering of the above questions before it can be completely evaluated by the community.