TULUNGAN: A SLANDERING-RESISTANT REPUTATION SYSTEM FOR COLLABORATIVE WEB FILTERING SYSTEMS

Alexis V. Pantola[1], Susan Pancho-Festin[2] and Florante Salvador[3]

[1]De La Salle University, 2401 Taft Avenue, Manila, Philippines
pong.pantola@delasalle.ph
[2]University of the Philippines, Diliman, Quezon City, Metro Manila, Philippines
susan.pancho@up.edu.ph
[2] De La Salle University, 2401 Taft Avenue, Manila, Philippines
bong.salvador@delasalle.ph

ABSTRACT

Reputation systems measure the credibility of contributions and contributors in collaborative web systems. Measuring the credibility is significant since a collaborative environment generally allows anyone with Internet access to provide contribution.

Collaborative web systems are susceptible to malicious users who intentionally provide inaccurate contents. With the help of reputation systems, the effect of such malicious activities can be reduced.
Reputation systems allow Internet users to rate the contributions made by other users. However, there are malicious users who will go beyond providing wrong contributions. They attempt to make reputation systems useless by launching attacks such as slandering. Slandering happens when a malicious user or group of malicious users intentionally provide a negative rating to accurate contributions provided by good users. Such activity lowers the reputation of good users and in most cases it even helps improve the reputation of slandering users.

This paper presents a reputation system called Tulungan that is designed to measure the contributor and rater reputation of users of a collaborative web system that is used for web filtering. User contributions are in the form of URL categorizations. It is the role of Tulungan to determine the correctness of the categorizations. A simulation is presented to validate the resilience of Tulungan in the presence of slandering users. The result of the simulation shows that Tulungan is not only resistant to slandering but it is still effective even if the number of good users is less than its slandering counterpart. Even if there are only 15% good users, the number of correct URL categorizations outnumbers incorrect contributions.

KEYWORDS

Reputation System, Collaborative Web Systems, Web Filtering

1. INTRODUCTION

Collaborative web systems such as Untangle [1] and Wikipedia [2] allow anyone with Internet access to contribute information (e.g., website categories in Untangle and articles in Wikipedia). Such approach poses a problem since malicious contributors will deliberately provide inaccurate contributions, thus, making these collaborative web systems useless or unreliable.
Reputation systems address problems of collaborative web systems by measuring the credibility of contributors as well as their contributions [3].

Collaborative web systems that employ a reputation system allow users to assume the roles of contributors and raters. Contributors provide contents to collaborative web systems, while raters assess the correctness of the contributed contents.

Reputation systems detect inaccurate contents that may come from malicious contributors through the assessment made by raters. In some cases, reputation systems may also promote cooperation [4]. As a result, contributors are not only discouraged from performing malicious behaviours, they are also encouraged to regularly contribute accurate contents.

Reputation systems are generally effective if raters provide accurate assessment. However, just like contributors, there may be malicious raters who deliberately provide incorrect ratings (i.e., giving negative ratings to accurate contributions or positive ratings to inaccurate contributions). One way to solve this is to limit the role of raters to a set of content managers. This approach is performed by web systems such as Untangle [5]. The content managers are responsible for sustaining and nurturing the contributions [6]. However, this makes the rating process unscalable since the verification process is done by a fewer number of people relative to contributors.

Allowing anyone to rate (i.e., not limiting the rating process to content managers) is more scalable. However, as mentioned earlier, it is prone to raters who intentionally provide incorrect assessment. Fortunately, most reputation systems are still effective even if there are malicious raters. However, they are consensus-dependent and assume that the number of malicious raters are less than their good counterpart (i.e., those that provide correct rating) [7]. Once the good raters are outnumbered by malicious ones, the reputation system fails to measure correctly the credibility of contributors (i.e., good raters are given lower reputation values relative to their malicious counterpart) [8]. There are even cases when malicious raters will employ attacks on reputation systems such as slandering in order to magnify the bad effect they provide to collaborative web systems.

Slandering happens when a malicious rater intentionally provides a negative rating to a correct contribution made by a specific good contributor. His objective is to bring down the reputation of a good contributor. Motivations of such act may be revenge or intimidation. There are even cases when a malicious rater will resort to Sybil attack to further increase the effect of slandering [9]. Sybil attack happens when a user creates phantom accounts and use these to give negative rating to a good contributor.

There are reputation systems that discourage slandering. However, they are designed to work in mobile ad-hoc networks (MANETs) [10] to encourage nodes in MANETs to forward messages of other nodes. In addition, their solution to prevent slandering may not be practical for collaborative web systems. For example, OCEAN [11] and LARS [12] discourages slandering by not allowing negative feedback. However, negative feedback may not be prevented in collaborative web systems since it limits the freedom of raters in giving negative assessment to inaccurate contents.

In this paper, a user that both contributes and rates maliciously is referred to as a malicious user. More specifically, a malicious user who rates correct contributions negatively is known as a slandering user.

2. REVIEW OF RELATED WORK

Reputation systems can be classified into three, namely content-driven and user-driven reputation systems, and reputation systems for MANETs.

2.1. Content-Driven Reputation System

A content-driven reputation system depends on the content of a contribution to determine its correctness. For example, WikiTrust [13, 14, 15] relies on the contents of Wikipedia pages to determine their accuracy. The stability of a content is equated to its credibility. It assumes that the less frequent an entry of a Wikipedia page is changed, the more credible it is since reviewers find it as already accurate and does not require any correction. However, such an assumption may lead to a wrong conclusion since a "non-edit" to an entry does not necessarily imply that said entry is correct. There are cases when authors are "lazy" in reviewing and editing an entire article and focus only in entries that interest them thereby leaving other entries unedited. For example, a biography entry in Wikipedia is proven inaccurate even if it is not edited for 132 days [16]. If the credibility of this entry is measured using WikiTrust, it is possible that it will be incorrectly labelled as accurate.

2.2. User-Driven Reputation System

A user-driven reputation system relies on user rating to measure the accuracy of contents. They are highly utilized in websites that allow Internet users to contribute contents and rate these contributions.

They may vary in names (e.g., rating system of Epinions [17], customer feedback and rating platform of BizRate [18], feedback forum of eBay [19], karma system of Reddit [20], social recommendation of Digg [21], moderation system of Slashdot [22], Rater-rating reputation system [23]), but all of them have the objective of assessing the accuracy of contributions in collaborative web systems.

User-driven reputation system assumes the cooperation among authors to become effective. As an example, eBay uses a Feedback Forum [24] in order to allow users to provide positive, neutral, or negative ratings [25] every time a transaction is completed. However, expecting an accurate feedback may be a challenge since users may not give a negative feedback for fear of retaliation or intentionally provide negative feedback just to cause slandering [26].

Digg's social recommendation also suffer a similar problem since a study suggests that its initial recommendation process can be influenced by the social neighbourhood of contributors. This means that contributions may be recommended due to through network effect and not due to their quality [27].

Slashdot's moderation system solves the problem of other user-driven reputation systems by employing meta-moderators [28, 29]. However, it has an issue on scaling since the ratio of meta-moderators with respect to contributors may be very large and prohibits meta-moderators to review all incoming contributions.

2.3. Reputation Systems in MANETs

There are several reputation systems in MANETs that can be adopted to discourage slandering. Watchdog and Pathrater (WP) [30], CORE [31], CONFIDANT [32], OCEAN [11], and LARS

[12] provide mechanism that discourages nodes from reporting false misbehaviour in a routing protocol called DSR [33]. However, such mechanisms limit the capability of the reputation systems. As an example, OCEAN and LARS are limited to positive feedback in order to avoid slandering.

2.4. Comparison of Algorithms

Table 1. Comparison of Different Systems.

System	for Web Systems	User-Driven	needs a Content Manager	Resistant to Slandering	Consensus-Independent
WikiTrust	yes	no	no	yes	yes
Rating System of Epinions	yes	yes	no	no	no
BizRate Customer Feedback and Ratings Platform	yes	yes	no	no	no
Feedback Forum of eBay	yes	yes	no	no	no
Karma System of Reddit	yes	yes	no	no	no
Digg's Social Recommendation	yes	yes	no	no	no
Slashdot's Moderation System	yes	yes	yes	yes	yes
Rater Rating	yes	yes	no	yes	no
Watchdog and Pathrater	no	no	no	no	no
CONFIDANT	no	no	no	no	no
CORE	no	no	no	yes	no
OCEAN	no	no	no	yes	no
LARS	no	no	no	yes	no
Target Characteristics	yes	yes	no	yes	yes

Table 1 provides a comparison of the different reputations systems discussed in this section. It is based from the study presented in [34]. The target characteristics of the reputation system considered in this study are shown in the last row. For easier comparison, table cells that are shaded are those characteristics that match the target.

WikiTrust satisfies all the targeted characteristics except it is a content-driven reputation system. The rating system of Epinions, BizRate customer feedback and ratings platform, feedback forum of eBay, karma system of Reddit, and Digg's social recommendation are all susceptible to slandering and are consensus-dependent.

Epinion employs Eroyalties credits to encourage users to provide high quality contribution. However, such credits do not prevent malicious users from performing slandering in order to deprive good contributors from getting Eroyalties credits.

The feedback forum of eBay accepts only feedback from users who were involved in a particular transaction (i.e., seller or buyer). This reduces but does not eliminate the chances of slandering since Sybil attack cannot be performed.

The moderation system of Slashdot satisfies all the targeted characteristics except for its dependence on a content manager. This makes it relatively not scalable compared to reputation systems that do not employ the services of content managers.

Similar to Slashdot's moderation system, Rater Rating also missed on a single targeted characteristic. It is not consensus-independent. However, just like the moderation system of Slashdot, it is not susceptible to slandering.

Among the five reputation systems for MANETs, CORE, OCEAN, and LARS are not susceptible to slandering. However, they mitigate this problem by disallowing or ignoring negative feedback which may not be applicable for most collaborative web systems.

3. THE TULUNGAN REPUTATION SYSTEM

This section presents a summary of the algorithm used by Tulungan as presented in [7].
The algorithm is divided into three phases: initialization, contribution and rating, and computation (refer to Algorithm 1). The initialization phase happens every time a new user and/or URL category are introduced in the reputation system. The last two phases are repeated every month.

Algorithm 1. Tulungan Reputation System.

Initialization Phase:

1. Initialize the contribution reputation φ_c, rating reputation φ_r, and level α of all URL categories s in $\mathbf{S}$.
2. Initialize the contributor reputation ρ_c and rater reputation ρ_r of user u and add it in $\mathbf{U}$.

Contribution and Rating Phase:

3. Allow all users u to add contribution c in $\mathbf{C}$.
4. Determine potential raters p and add them in $\mathbf{P}$.
5. Allow all users u that are potential raters to add rating group g in $\mathbf{G}$ and rating r in $\mathbf{R}$.

Computation Phase:

6. Update the rating reputation φ_r of all URL categories s in $\mathbf{S}$.
7. Update the contribution reputation φ_c of all URL categories s in $\mathbf{S}$.
8. Update the level α of all URL categories s in $\mathbf{S}$.
9. Update the contributor reputation ρ_c of all users u in $\mathbf{U}$.
10. Update the rater reputation ρ_r of all users u in $\mathbf{U}$.
11. Update the overall reputation ρ_o of all users u in $\mathbf{U}$.

The initialization phase sets the contribution reputation and rating reputation values of all URL categories as well as the contributor and rater reputation values of users. The contributor and rater reputation values will start with a value close to zero. As the users provide good contributions and ratings, their reputation values increase.

The contribution and rating phase allows users to provide contribution. Contribution is in the form of URL category. For example, a user can contribute that www.tennis.com is a sports site. In addition he/she can also say that www.tennis.com is NOT a pornographic site.

Aside from providing contributions, the potential raters are determined in the second phase. They are called potential raters since it is possible that users will fail to rate the contributions they are

requested to review and rate. Determining potential raters greatly depends on the contribution reputation of the URL categories as well as the contributor and rater reputation of users. Since there can be many contributions, this phase selects contributions that are worth reviewing (i.e., many users contributed the same URL categories).

Potential raters will be provided with a set (or several sets) of contributions to rate. Each set is composed of three URLs. The three URLs are composed of one unknown URL and two control URLs. The unknown URL is the one the Tulungan algorithm is interested in verifying if the contributions related to the URL is accurate. On the other hand, the categories of the two control URLs are already known by Tulungan. However, the two control URLs as well as the unknown URL are presented to potential raters in a way that the three are treated as unknown URLs. When a user rates the three URLs, Tulungan can check if the rating provided in the two control URLs are correct. If they are correct, Tulungan assumes that the user is ``serious'' in providing rating and therefore will also assume that the rating he/she gave on the unknown URL is also correct. This approach is similar to reCAPTCHA [35].

The use of control URLs helps in reducing the effect of slandering users. Since the control URLS verify the seriousness of rating, users who attempt to perform slandering will be detected by Tulungan.

The computation phase calculates the new contribution and rating reputation of URL categories as well as the contributor and rater reputation values of users. This is essential to measure the credibility of users in providing contributions as well as the accuracy of URL categories. The reputation values that were computed in this phase.

The contributor and rater reputation of users may be used as a basis for a reward system in Tulungan. Since Tulungan can be used in a collaborative web filtering system, users can receive category updates such that if they are utilizing a web filtering system, the filtering is based on an updated category list. However, Tulungan can limit the updates users receive based on their current contributor and rater reputation (i.e., low reputation means no update). With this approach, users are discouraged to misbehave (e.g., slander) to have high reputation values.

4. EVALUATION OF THE REPUTATION SYSTEM

4.1. Evaluation Set-up

The system is evaluated via simulation. Good, malicious and slandering users are modelled in order to verify their effects in reputation systems. Take note that the focus of the study is the evaluation of Tulungan in the presence of slandering users. However, the effect of malicious users are also presented in order to compare the effects of slandering users with malicious users.
A good contributor provides correct categorization of URLs most of the time. In the simulation, there is only 1 in 5000 chances a good contributor provides a wrong contribution. A good rater has a similar behaviour as the good contributor.

A malicious contributor provides incorrect categorization of URLs most of the time. In order to ensure such behaviour, the simulation is configured such that malicious contributors only has a 1 in 5000 chances of providing a correct contribution. Unlike good contributors and raters, a malicious rater is modelled differently from its contributor counterpart. Since the rating process requires rating the URL categories of three URLs (i.e., two control URLs and an unknown URL), a malicious rater needs to rate correctly the two control URLs and intentionally make an incorrect rating for the unknown. Giving a correct rating for the control URLs is essential, otherwise, the

reputation system can detect that the rater is giving a wrong rating and may be doing something malicious. However, since the malicious rater has no idea which of the three URLs are the controls and the unknown, he needs to make a guess. There is a 1 in 3 chances that a malicious rater will correctly guess which is the unknown URL. Therefore, it also has a 1 in 3 chances of being successful in giving a wrong rating for the unknown URL while remaining undetected by the reputation system.

Slandering is simulated by combining it with Sybil attack. The slandering users are formed as a group of 16. This group is further divided into two subgroups: A and B. Members of subgroups A and B contribute in the same way malicious contributors provide contribution. However, since Sybil attack is adopted, members of subgroup A provide the same set of contributions. Since most reputation systems do not allow users to rate their own contribution, it is impossible for members of subgroup A to rate their own contribution. This is the purpose of subgroup B. Members of subgroup B have a chance of rating the contributions of subgroup A. Whenever they are allowed to rate contributions of subgroup A, they ensure that they are given a positive rating such that the wrong contributions made by subgroup A will be considered correct. To perform slandering, the group of slandering users identify a good user to victimize and ensure that the contributions and ratings they make contradicts the contributions of their victim. For example, if a good user categorizes www.tennis.com as a sports website, subgroup A categorizes www.tennis.com as a non-sports website. In addition, subgroup B rates positively the contributions of subgroup A.

To keep track of the user type (e.g., a good contributor and rater), the simulation uses two fields for each user: contributor_type and rater_type. This is essential to determine the contribution and rating behaviour that a user will perform in the simulation. However, these fields are not used in the actual computation of the contributor, rater, and overall reputation of each user.
In addition, the number of correct and wrong URL categorizations are also recorded in order to verify the effectiveness of the reputation systems in decreasing the occurrence of wrong contributions.

The simulation does not only cover the Tulungan reputation system but as well as an ordinary reputation system in order to compare the two systems in terms of differentiating good, malicious, and slandering users based on their reputation calculation.

An ordinary reputation system has the following characteristics:

- similar to Tulungan, it does not allow users to rate their own contribution
- unlike Tulungan, users are allowed to choose contributions to rate (as long as they are not their own contributions)
- unlike Tulungan, it does not employ control URLs during rating, which makes it harder to detect malicious ratings

The simulation is divided into four (4) parts as specified in Table 2.

Table 2. Simulation Experiments.

Type of Users Involved	Ordinary	Tulungan
Good vs. Malicious	Simulation 1	Simulation 2
Good vs. Slandering	Simulation 3	Simulation 4

Each simulation is executed with a varying percentage of good users relative to their malicious and slandering counterpart. For example, in simulations 1 and 2, they start with 5% good users

and 95% malicious users. They are executed again with 10% good and 90% malicious users. This is repeated with an increment of 5% in the number of good users until the percentage of good users reach 95%. Take note that even if a slandering user is a specific type of malicious user, simulation results pertaining to malicious users cover only non-slandering users.

For each execution, the simulation is composed of 500 users and runs for 366 simulated days (i.e., January 1 to December 31, 2000). Everyday, each user in the simulation provides one contribution. On the 1st day of each month, the potential raters are determined by the reputation system. On the 2nd day of each month, each user provides a rating based on the potential rater list generated on the 1st day. On the 28th day of each month, the contributor reputation and rater reputation of all the users are determined. At the end of the 366th day, the average of the contributor reputation, rater reputation, and overall reputation of all good users is computed. The same thing is done with the reputation of malicious and slandering users. To get the average reputation for each user type, the fields contributor_type and rater_type are used. Take note that fields are used in getting the average reputation and not in the computation of individual reputation of the users. In addition, the number of correct and wrong URL contributions are counted by comparing the URL categorizations determined through the contributions (and verified with ratings), versus the actual URL categorizations.

4.2. Results and Analysis

In order to verify the consistency of the experiments, each simulation is run four times. The results presented in this paper are derived from the first run of the simulation. The results of the second to fourth runs are generally similar to the first.

4.2.1 Good vs. Malicious Users (Simulations 1 and 2)

Figures 1 and 2 are the results of simulations 1 and 2, respectively.

As shown in Figure 1a, the ordinary reputation system gives good users a higher contributor reputation when there is at least 50% good users. The same can be observed in the rater and overall reputation of the ordinary reputation system as seen in Figures 1b and 1c. The results confirm the consensus-dependence of the ordinary reputation system.

Similar to the reputation of good users, the number of correct URL categorizations surpasses the number of wrong ones when there are 50% good users (refer to Figure 1d). This is a critical problem with consensus-dependent reputation systems since a significant number of wrong URL categorizations are produced even if the number of malicious users is just slightly higher than their good counterpart. For example, Figure 1d shows that with 55% malicious users (or 45% good users) there are approximately 4000 wrong URL categorizations and almost no correct URL categorization.

Figure 2a shows the contributor reputation of users as computed by Tulungan. Even if there are only 20% good users, their contributor reputation is higher than malicious users.

The rater reputation computed by Tulungan is illustrated in Figure 2b. As shown in the graph, even with only 5% good users, malicious users are outperformed by their good counterpart. It should be noted that the rater reputation of good users is below the mid-range of 5. However, their rater reputation goes beyond 5 when there are at least 20% good users.

Figure 2c shows that Tulungan is able to limit the overall reputation of malicious users to approximately 0, while good users have at least an overall reputation of greater than 5 at the 20% mark.

It is illustrated in Figure 2 that Tulungan is a consensus-independent reputation system. It requires only 20% good users to become effective. In addition, Figure 2d shows that the number of correct URL categorizations is more than its wrong counterpart even if there are only 20% good users. More importantly, even if there are less than 20% good users, the number of wrong URL categorizations is less than 200. This is significantly less than the wrong URL categorizations produced in the ordinary reputation system when there are only 5% good users. As shown in Figure 1d, wrong URL categorizations can reach more than 5000.

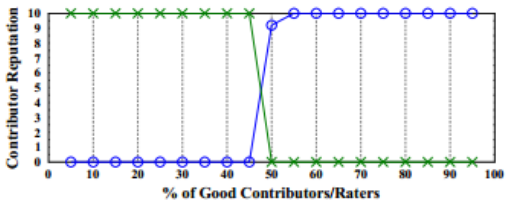

(a) Contributor Reputation using an Ordinary Rep. Sys.

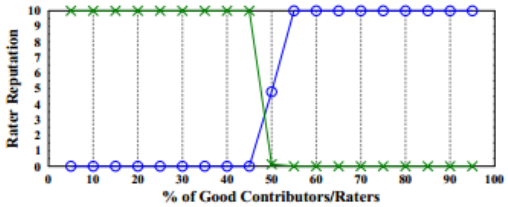

(b) Rater Reputation using an Ordinary Rep. Sys.

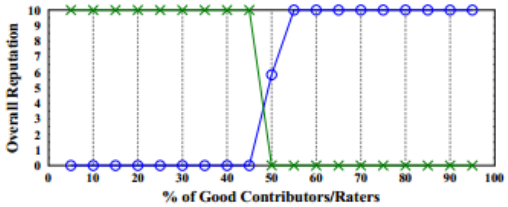

(c) Overall Reputation using an Ordinary Rep. Sys.

User Legend: ○ - good × - malicious

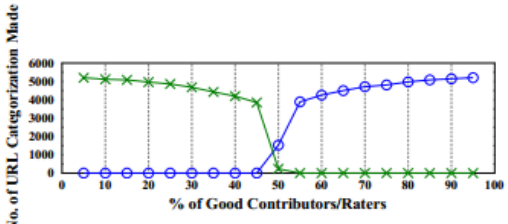

(d) Categorization Count using an Ordinary Rep. Sys.

Categorization Legend: ○ - correct × - wrong

Figure 1. Good versus Malicious Users: Ordinary Reputation System Results (Simulation 1)

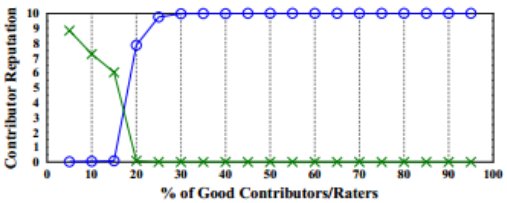

(a) Contributor Reputation using Tulungan

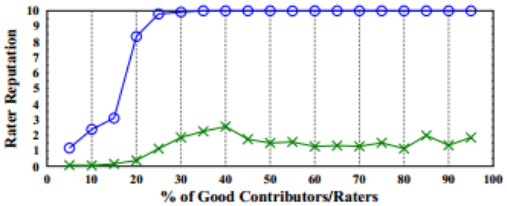

(b) Rater Reputation using Tulungan

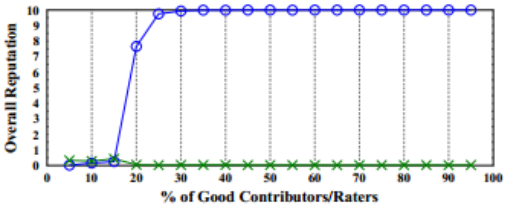

(c) Overall Reputation using Tulungan

User Legend: ○ - good × - malicious

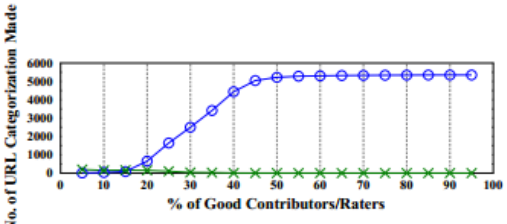

(d) Categorization Count using Tulungan

Categorization Legend: ○ - correct × - wrong

Figure 2. Good versus Malicious Users: Tulungan Results (Simulation 2)

4.2.1 Good vs. Slandering Users (Simulations 3 and 4)

Figures 3 and 4 are the results of simulations 3 and 4, respectively.

As shown in Figures 3a, 3b, and 3c, even if there are more good users relative to slandering users (i.e., 75% good users), the latter can outperform the former in terms of contributor, rater, and overall reputation. Furthermore, only 25% slandering users is needed such that correct URL categorizations are outnumbered by wrong ones. This is shown in Figure 3d.

Figures 4a, 4b, and 4c show that Tulungan is very resilient to slandering. It requires only 15% good users such that their contributor reputation is higher than their slandering counterpart. In addition, the rater reputation of good users is higher than slandering users even if there are 95% slandering users. With Tulungan, the overall reputation of slandering users is prevented from going above 0.5 regardless of the number of good users.

Similar to simulation 2 (good versus malicious), the number of wrong URL categorizations is less than 200 regardless of the number of good users. On the other hand, correct URL categorizations can reach more than 5000.

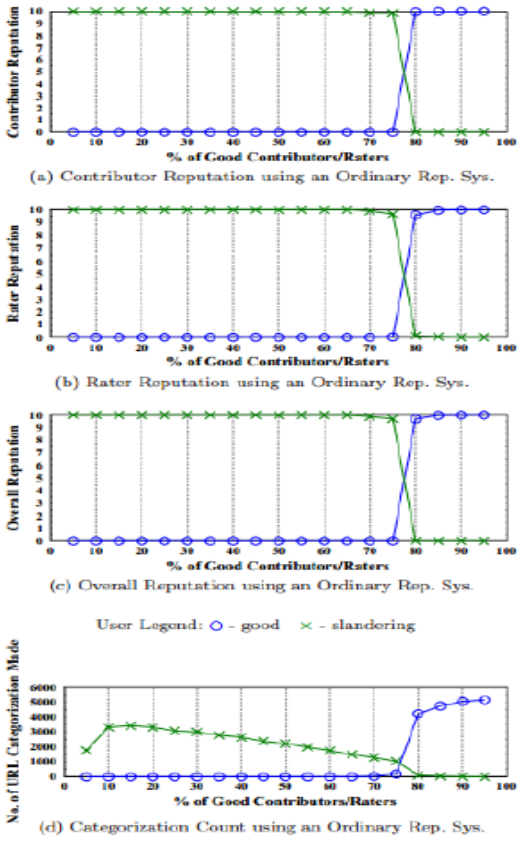

Figure 3. Good versus Slandering Users: Ordinary Reputation System Results (Simulation 3)

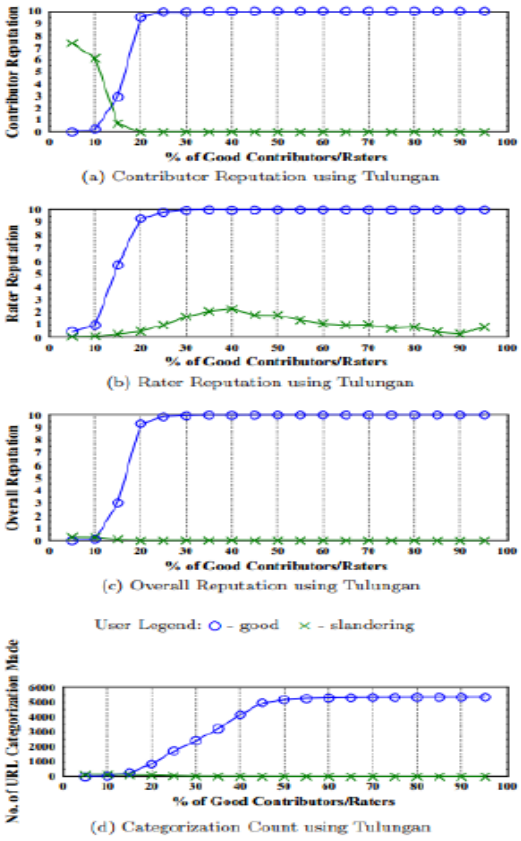

Figure 4. Good versus Slandering Users: Tulungan Results (Simulation 4)

Table 3 shows a summary of the required percentage of good users for each reputation system when there are malicious and slandering users.

Tables 4 and 5 summarizes the maximum number of correct and wrong URL categorizations that were determined in simulations 1 to 4.

Table 3. Percentage of Good Users needed to make a Reputation System Effective.

Type of Users Involved	Ordinary	Tulungan
Good vs. Malicious	50%	20%
Good vs. Slandering	80%	15%

Table 4. Maximum Number of Correct and Wrong URL Categorizations in Simulations 1 and 2 (Good vs. Malicious Users).

URL Categorization	Ordinary	Tulungan
Correct	5202	5349
Wrong	5199	195

Table 5. Maximum Number of Correct and Wrong URL Categorizations in Simulations 3 and 4 (Good vs. Slandering Users).

URL Categorization	Ordinary	Tulungan
Correct	5147	5348
Wrong	3414	140

5. CONCLUSION

The simulation results show that Tulungan is slandering-resistant. Tulungan requires only 15% good users in order to become effective. In contrast, the ordinary reputation system is effective only when there is at least 50% good users. In addition, regardless of the number of slandering users, the maximum number of wrong URL categorizations is less than 3% (or 140/5348)the maximum number of good URL categorizations that can be achieved. In the simulation, slandering users can only produce less than 200 wrong URL categorizations while good users can produce more than 5000 good ones when Tulungan is used as the reputation system. On the other hand, the ordinary reputation system produces more than 3000 wrong URL categorizations or more than 66% (or 3414/5147) wrong URL categorizations relative to the maximum number of correct ones.

Unlike the various reputation systems specified in Table 1, Tulungan is able to satisfy all the enumerated target characteristics of a reputation system. Tulungan can be used to effectively categorize the URLs of websites even if there are more malicious users relative to their good counterpart. Furthermore, Tulungan is still effective in providing categorization of URLs even if slandering users are present and Sybil attack is performed.

Although Tulungan is slandering-resistant, further study can be performed in order to assess its effectiveness against other reputation attacks such as whitewashing, orchestration, and denial of service.

REFERENCES

[1] Untangle - Multi-functional firewall Software -Open Source Content Filter and Spam Filter [On-line]. URL http://www.untangle.com. Last accessed March 2011

[2] Wikipedia, the free Encyclopedia [Online]. URL http://www.wikipedia.org. Last accessed January 2011

[3] J. Kennes, A. Schiff, The Value of a Reputation System. Industrial Organization 0301011, Econ-WPA (2003). URL http://ideas.repec.org/p/wpa/ wuwpio/0301011.html

[4] L. Buttyan, J.P. Hubaux, Mob. Netw. Appl. 8, 579 (2003). DOI http://dx.doi.org/10.1023/A:1025146013151. URL http://dx.doi.org/10.1023/A:1025146013151

[5] Untangle - Web Filter submission tool [On-line]. URL http://forums.untangle.com/web-filter/2143-web-filter-submission-tool.html

[6] C. Gray. Launching a Collaboration or Content Management System: 8 Tricks for Adoption [On-line] (2010)

[7] A.V. Pantola, S. Pancho-Festin, F. Salvador, Science Diliman 23(2) (2012)

[8] E. Yao, R. Fang, B.R. Dineen, X. Yao, Journal of Business Research 62(12), 1281 (2009)

[9] J.R. Douceur, In Proceedings for the 1st International Workshop on Peer-to-Peer Systems (IPTPS) pp. 251–260 (2002)

[10] IETF MANET Working Group. Mobile Ad Hoc Networks (MANET). WorkingGroup charter [On-line]. URL http://www.ietf.org/html.charters/manet-charter.html

[11] S. Bansal, M. Baker. Observation-based Cooperation Enforcement in Ad hoc Networks (2003). URL http://arxiv.org/abs/cs/0307012

[12] J. Hu, M. Burmester, in Proceedings of the 44th annual Southeast regional conference (ACM, New York, NY, USA, 2006), ACM-SE 44, pp. 119–123. DOI http://doi.acm.org/10.1145/1185448.1185475. URL http://doi.acm.org/10.1145/1185448.1185475

[13] B.T. Adler, L. de Alfaro, in Proceedings of the 16th International Conference on World Wide Web (ACM, New York, NY, USA, 2007), WWW '07, pp. 261–270. DOI http://doi.acm.org/10.1145/1242572.1242608. URL http://doi.acm.org/ 10.1145/1242572.1242608

[14] B.T. Adler, K. Chatterjee, L. de Alfaro, M. Faella, I. Pye, V. Raman, in Proceedings of the 4th International Symposium on Wikis (ACM, New York, NY, USA, 2008), WikiSym '08, pp. 26:1–26:12. DOI http://doi.acm.org/10.1145/1822258.1822293. URL http://doi.acm.org/10.1145/1822258.1822293

[15] B.T. Adler, L. de Alfaro, I. Pye, V. Raman, in Proceedings of the 4th International Symposium on Wikis (ACM, New York, NY, USA, 2008), Wik-iSym '08, pp. 15:1–15:10. DOI http://doi.acm.org/10.1145/1822258.1822279. URL http://doi.acm.org/10.1145/1822258.1822279

[16] T. Cross, First Monday 11(9) (2006). URL http://firstmonday.org/issues/issue11_9/cross/index.html

[17] Epinions.com - FAQs [Online]. URL http://www99.epinions.com/help/faq/. Last accessed June 2011

[18] BizRate Retailer Ratings [Online]. URL http://merchant.shopzilla.com/oa/customer_ratings/

[19] Feedback Forum - How Feedback works [On-line]. URL http://pages.ebay.ph/services/forum/feedback.html. Last accessed March 2011

[20] Reddit.com: Help [Online]. URL http://www.reddit.com/help/faq. Last accessed June 2011

[21] Digg [Online] (2012). URL http://www.digg.com

[22] Slashdot - FAQ [Online]. URL http://slashdot.org/faq. Last accessed June 2011

[23] A.V. Pantola, S. Pancho-Festin, F. Salvador,in Proceedings of the 3rd international conference on Security of information and networks (ACM, New York, NY, USA, 2010), SIN '10, pp. 71–80. DOI http://doi.acm.org/10.1145/1854099.1854116. URL http://doi.acm.org/10.1145/1854099.1854116

[24] P. Resnick, K. Kuwabara, R. Zeckhauser, E. Friedman, Commun. ACM 43, 45 (2000). DOI http://doi.acm.org/10.1145/355112.355122. URL http://doi.acm.org/10.1145/355112.355122

[25] Feedback - 5 Simple Steps to eBay Feedback [On-line]. URL http://pages.ebay.ph/ebayexplained/feedback.html. Last accessed March 2011

[26] D. Steiner, Auction Bytes. The Independent Trade Publication for Online Merchants (2003). URL http://www.auctionbytes.com/cab/abu/y203/m01/abu0087/s02

[27] K. Lerman, A. Galstyan, in Proceedings of the first workshop on Online social networks (ACM, New York, NY, USA, 2008), WOSP '08, pp. 7–12. DOI http://doi.acm.org/10.1145/1397735.1397738. URL http://doi.acm.org/10.1145/1397735.1397738

[28] C. Lampe, Ratings Use in an Online Discussion System: The Slashdot Case. Ph.D. thesis, University of Michigan (2006)

[29] C. Lampe, P. Resnick, in Proceedings of the SIGCHI conference on Human factors in computing systems (ACM, New York, NY, USA, 2004), CHI '04, pp. 543–550. DOI http://doi.acm.org/10.1145/985692.985761. URL http://doi.acm.org/10.1145/985692.985761

[30] S. Marti, T.J. Giuli, K. Lai, M. Baker, in Proceedings of the 6th Annual International Conference on Mobile Computing and Networking (ACM, New York, NY, USA, 2000), MobiCom '00, pp.255–265. DOI http://doi.acm.org/10.1145/345910.345955. URL http://doi.acm.org/10.1145/345910.345955

[31] P. Michiardi, R. Molva, in Proceedings of the IFIP Communication and Multimedia Security Conference (2002)

[32] S. Buchegger, J.Y. Le Boudec, in Proceedings of the Tenth EuromicroWorkshop on Parallel, Distributed and Network-based Processing (2002), pp. 403–410

[33] D.B. Johnson, D.A. Maltz, J. Broch, in Ad Hoc Networking, edited by Charles E. Perkins, Chapter 5 (Addison-Wesley, 2001), pp. 139–172

[34] K. Hoffman, D. Zage, C. Nita-Rotaru, ACM Comput. Surv. 42, 1:1 (2009). DOI http://doi.acm.org/10.1145/1592451.1592452. URL http://doi.acm.org/10.1145/1592451.1592452

[35] reCAPTCHA - Stop Spam, Read Books[Online]. URL http://www.google.com/recaptcha/. Last accessed March 2011

Authors

Alexis V. Pantola is a former faculty member of the Computer Technology Department at De La Salle University – Manila. He finished his Masters in Computer Science degree in De La Salle University. He specializes in computer networks, information security, and data analytics.

Susan Pancho-Festin is head of the Computer Security Group at the University of the Philippines-Diliman. She obtained her PhD from Cambridge University, where she was part of the Security Group at the Computer Laboratory. Prior to this, she graduated from Royal Holloway University of London with an MSc in Information Security degree. Her research interests are in security protocols and secure software engineering.

Florante R. Salvador is a full-time faculty member at the College of Computer Studies, De La Salle University, Manila. His current research interest is on computer graphics. He received his Dr. Eng. degree fromYamanashi University, Japan in 1996.