Resistance of Trust Management Systems Against Malicious Collectives

Miroslav Novotný and Filip Zavoral(✉)

Faculty of Mathematics and Physics, Charles University in Prague,
118 00 Prague, Czech Republic
{novotny, zavoral}@ksi.mff.cuni.cz

Abstract. Malicious peers in Peer-to-peer networks can develop sophisticated strategies to bypass existing security mechanisms. The effectiveness of contemporary trust management systems is usually tested only against simple malicious strategies. In this paper, we propose a simulation framework for evaluation of resistance of trust management systems against different malicious strategies. We present results of five TMS that represent main contemporary approaches; the results indicate that most of the traditional trust managements are vulnerable to sophisticated malicious strategies.

Keywords: Trust management · Peer to peer networks

1 Introduction

One of the promising architectures of large-scale distributed systems is based on peer to peer architecture (P2P). However, providing proper protection to such systems is tricky. The P2P applications have to deal with treacherous peers that try to deliberately subvert their operation. The peers have to trust the remote party to work correctly. The process of getting this trust is, however, far from trivial.

Many trust management systems (TMS) have been developed to deal with treacherous peers in P2P networks. The main idea of these systems is sharing experience between honest peers and building reputations. Nevertheless, the group of cooperating malicious peers is often able to bypass their security mechanisms and cause a great deal of harm. The malicious collectives represent the main reasons why managing trust represents the biggest challenge in the current P2P networks.

In this paper, we investigate several TMSs and use the simple taxonomy to organize their major approaches. Using the simulation framework called P2PTrustSim we investigate different strategies used by malicious peers. Beside traditional strategies, we propose new, more sophisticated strategies and test them against five trust management systems. These systems have been chosen as the representatives of major approaches. Our goal was to verify the effectiveness of various TMSs under sophisticated malicious strategies. We have chosen five contemporary TMS: EigenTrust [1], PeerTrust [2], H-Trust [3], WTR [4], and BubbleTrust [5]. These TMS represent main contemporary approaches in Trust Management.

P.C. Vinh et al. (Eds.): ICCASA 2013, LNICST 128, pp. 67–76, 2014.
DOI: 10.1007/978-3-319-05939-6_7, © Springer International Publishing Switzerland 2014

2 Malicious Strategies and Evaluation Criteria

In order to facilitate comparison of different TMSs and their behaviour under different malicious strategies we created a simulation framework [6] called P2PTrustSim. We used FreePastry, a modular, open-source implementation of the Pastry [12], P2P structured overlay network. Above the FreePastry, we created the peer simulation layer which implements various peers' behaviour.

2.1 Malicious Strategies

Most of the TMSs work well against straightforward malicious activities. However, the malicious peers can develop strategies to maintain their malicious business. Each peer can operate individually but the biggest threat is the collusion of malicious peers working together.

2.1.1 Individual Malicious Strategies

These strategies do not involve the cooperation between malicious peers.

False Meta-data - Malicious peers can insert false, attractive information into the meta-data describing their bogus resources to increase the demands for them.

Camouflage - The malicious peers that are aware of the presence of the TMS can provide a few honest resources. There can be many variants of this strategy, differing in the ratio of honest and bogus services or the period between changing behaviour. In some literature, the variant of this strategy is called Traitors [7, 8–10].

2.1.2 Collective Malicious Strategies

Malicious peers have a significantly higher chance to succeed if they work in a cooperative manner; this is considered as the biggest treat for P2P applications [11].

Full Collusion - All members of a malicious collective provide bogus resources and create false positive recommendations to all other members of the collective.

Spies - The malicious collective is divided into two groups: spies and malicious. The spies provide honest services to earn a high reputation and simultaneously provide false positive recommendations to the malicious part of the collective.

2.1.3 Newly Proposed Malicious Strategies

We analyzed published TMSs and known malicious tactics carefully and we suggest three new collective malicious strategies. Each strategy is designed for a particular type of TMS and tries to exploit its specific weakness.

Evaluator Collusion - If the TMS assesses credibility of the feedback source according to the truthfulness of its previous feedback, malicious peers can try to trick the TMS by using the services from peers outside the collective and evaluating them correctly. This feedback increases the credibility of malicious peers as recommenders and gives more weight to their feedback towards other members of collective.

Evaluator Spies - This strategy is a combination of Evaluator Collusion and Spies. The spies implement three techniques to maintain a credibility as a feedback source: they provide honest service, they use resources outside the collective and evaluate them correctly, and they create positive recommendations towards other spies.

Malicious Spies - This slight modification of the previous strategy is based on the idea that spies do not require a high reputation as resource providers. They can provide bogus resources and generate negative recommendations between each other. These recommendations are still truthful and should increase their credibility.

2.2 Evaluation Criteria

Each transaction within the framework is categorized on both sides (provider and consumer). The categories distinguish the type of the peer (honest or malicious), on which side of the transaction the peer was (provider or consumer), and the result of the transaction. The ulterior transactions represent honest transactions which malicious peers have to perform to fix their reputation. If no provider is sufficiently trustful, the transaction is refused and counted as ConsumeRefused. The originated peer typically tries to pick different service and repeat the transaction.

Let us suppose that all the malicious peers cooperate within a malicious collective in the network and all transactions from honest peers are honest. Our primary goal is to evaluate the success of each malicious strategy in different TMSs. Therefore, we defined four criteria:

MaliciousSuccessRatio (MSR) is a ratio between bogus transactions provided by malicious peers in the network with TMS and in the network without TMS (DummyTrust). It reflects the contribution of the given TMS and it is defined by the following formula:

$$MaliciousSuccessRatio = \frac{TotalBogus_{withTMS}}{TotalBogus_{withoutTMS}}$$

BogusRatio (BR) is a ratio between bogus and all services consumed by the honest peers. It is defined by the following formula:

$$BogusRatio = \frac{100 * TotalBogus}{\sum ConsumeHonest + TotalBogus}$$

MaliciousCost (MC) monitors the load associated with a malicious strategy. It is a ratio between extra transactions performed by the malicious peers to trick the TMS and the bogus transactions in the network. These extra transactions include faked and ulterior transactions and represent additional overhead for malicious peers which they try to minimize. We defined it by the following formula:

$$MaliciousCost = \frac{TotalUlterior + TotalFaked/2}{TotalBogus}$$

This metric gives us an idea of how much computational power and network utilization is required for a given malicious strategy.

The last criterion is a **MaliciousBenefit (MB)**. It represents how much beneficial transactions the malicious peers have to perform to pass one malicious service. It is defined by the following formula:

$$MaliciousBenefit = \frac{TotalUlterior}{TotalBogus}$$

The value above 1 means that there is benefit from the malicious collective which is bigger than the damage caused by the collective.

3 Simulation Results

We focused on two problems: the effectiveness of the strategies and the reaction of the TMSs to changes in peers' behaviours. The first problem was studied in a network that contains 200 peers and 80 peers are malicious; 40 % of nodes in the network are malicious, which represents a very dangerous environment. The honest peers wake up every 10 min and use one service from the network. The malicious peers also wake up every 10 min and perform a given number of faked or ulterior transactions. We ran 56 different simulations (7 TMSs each with 8 strategies). Each of the simulations represents 24 h. The data is counted in the last hour of the simulations when the TMSs are stabilized. Each simulation was repeated 20 times and average values are taken. The variation of results is expressed in the form of a relative standard deviation (RSD). The size of the network was designed with regards to simulation possibilities of the FreePastry and the load produced by our simulation. The results of other series of tests with the different settings were almost identical.

We set similar parameters for all TMSs. The most important parameter is the history period which determines how long the peers remember the information about previous transactions. We set this parameter to 30 cycles (5 h, in order to have a history period appropriate to the total simulation time) in all TMSs. The EigenTrust is not able to work correctly without pre-trusted peers, so we had to set 10 % honest peers as pre-trusted. Therefore, the EigenTrust has an advantage over other TMSs. Also, the numbers of ulterior and faked transactions are the same for all malicious strategies which use them.

3.1 Representative TMSs

The first simulations were performed in the network without TMS (DummyTrust) and in the network with the simplest version of TMS (SimpleTrust). We focused on the number of bogus transactions; these values will be used as a base for calculation of MaliciousSuccessRatio for other TMSs. The results are shown in Table 1.

As expected, the False Meta-data is the only useful strategy in DummyTrust. Other malicious mechanisms are useless or even counterproductive. The strategies Malicious Individual, Full Collusion, Evaluator Collusion and Malicious Spies have

Table 1. Number of bogus transactions in DummyTrust and SimpleTrust.

Strategy	DummyTrust		SimpleTrust		
	TotalBogus	RSD (%)	TotalBogus	RSD	Diff.
Simple Malicious Individual	262.20	8.95	247.15	6.25	6
Malicious Individual	435.65	3.05	387.70	4.12	11
Camouflage	310.85	4.02	281.20	3.54	10
Full Collusion	430.75	2.84	391.45	2.83	9
Evaluator Collusion	436.90	2.62	388.10	4.40	11
Spies	297.25	4.00	249.20	5.17	16
Evaluator Spies	301.30	4.31	244.80	7.00	19
Malicious Spies	433.65	2.54	386.65	3.17	11

almost the same results. All these strategies use False Meta-data, unlike The Simple Malicious Individual, which reaches fewer bogus transactions. The rest of the malicious strategies sacrifice a part of bogus transactions to circumvent TMSs, however these transactions have no effect in DummyTrust. The biggest variation in results has been measured in Simple Malicious Individual. In this strategy, honest peers completely rely on a random choice of communication partner.

The SimpleTrust has only slightly better results. The biggest improvement was measured in Evaluator Spies and Spies. These strategies are not suitable for simple TMSs. In fact, we have expected a bigger improvement. The limited factor is the size of the history period which was set to 30 cycles in all TMSs. Without cooperation with other peers, the information about peer's maliciousness is lost after 30 cycles and the delay between two transactions towards the same peer can be longer.

3.2 Efficiency Criterion

We measured the criteria described in Sect. 3. The most important of them is the MaliciousSuccessRatio (MSR); the measured values are in Tables 2 and 3 along with average numbers of bogus transactions and standard deviations. The MSR values above the threshold 0.5 are displayed in a bold font. We can see that only the BubbleTrust is resistant against all malicious strategies. There is at least one effective malicious strategy against all other TMSs. The EigenTrust, despite its advantage, is completely vulnerable to Spies and Evaluator Spies. These strategies are even able to perform more bogus transactions than it would be possible in a network without TMS. PeerTrust is resistant against only the simplest malicious strategies, on the other hand, malicious strategies like Evaluator Collusion and Evaluator Spies are 100 % effective. Also H-Trust does not work well, it is completely vulnerable to Evaluator Collusion and Evaluator Spies and the resistance against other strategies is not convincing either. WTR copes very well with individual strategies; especially the Camouflage is ineffective due to the risk factor. But the collective strategies can easily circumvent it. There are noticeable deviations in some malicious strategies. However, none of these deviations influence the MSR value that much that cross the limit 0.5.The next criterion is BogusRatio. Table 4 shows BogusRatio of each malicious strategy in all TMSs. In the worst case scenario, only 29 % of all transactions in the P2P network

Table 2. Malicious Success Ratio in BubbleTrust, EigenTrust and PeerTrust.

Strategy	BubbleTrust			EigenTrust			PeerTrust		
	Total Bogus	RSD (%)	MSR	Total Bogus	RSD (%)	MSR	Total Bogus	RSD (%)	MSR
Simple M Individual	6.2	66.2	0.0	64.2	21.6	0.2	17.5	20.3	0.1
Malicious Individual	1.5	89.9	0.0	137.9	15.2	0.3	0.0	0.0	0.0
Camouflage	1.55	84.96	0.00	87.85	24.82	0.28	200.60	6.84	0.65
Full Collusion	58.25	13.13	0.14	0.00	0.00	0.00	426.90	3.55	0.99
Evaluator Collusion	109.2	10.17	0.25	0.00	0.00	0.00	440.05	2.70	1.01
Spies	21.5	23.93	0.07	323.45	3.83	1.09	282.25	4.53	0.95
Evaluator Spies	48.3	11.44	0.16	295.50	28.08	0.98	300.00	4.65	1.00
Malicious Spies	53.5	11.21	0.12	0.55	–	0.00	297.95	4.87	0.69

Table 3. Malicious Success Ratio in HTrust and WTR.

Strategy	HTrust			WTR		
	Total Bogus	RSD (%)	MSR	Total Bogus	RSD (%)	MSR
Simple Malicious Individual	54.00	20.70	0.21	0.00	0.00	0.00
Malicious Individual	142.15	8.69	0.33	0.00	0.00	0.00
Camouflage	56.30	15.94	0.18	0.00	0.00	0.00
Full Collusion	138.05	8.23	0.32	435.45	2.26	1.01
Evaluator Collusion	411.00	4.30	0.94	436.70	3.73	1.00
Spies	108.10	8.74	0.36	293.30	5.57	0.99
Evaluator Spies	296.60	3.84	0.98	302.65	4.39	1.00
Malicious Spies	299.55	4.04	0.69	304.40	3.86	0.70

Table 4. BogusRatio for different malicious strategies and TMSs.

Strategy	EigenTrust (%)	H-Trust (%)	PeerTrust (%)	WTR (%)	BubbleTrust (%)
Simple M Individual	13	11	4	0	2
Malicious Individual	34	35	0	0	1
Camouflage	21	13	38	0	0
Full Collusion	0	34	72	73	18
Evaluator Collusion	0	70	73	73	29
Spies	55	22	48	49	5
Evaluator Spies	63	50	50	50	11
Malicious Spies	1	56	57	57	16

with the BubbleTrust can be bogus. Other TMS tolerate 63 % (EigenTrust), 70 % (H-Trust), 73 % (PeerTrust and WTR) bogus transactions.

Table 5 shows MaliciousCost of malicious strategies which use ulterior or faked transactions. Other strategies (Simple Malicious Individual and Malicious Individual) have no additional cost. MaliciousCost of the strategies with no measurable MSR is infinite and the cells contain ‘N/A’.

Table 5. MaliciousCost for different malicious strategies and TMSs.

Strategy	EigenTrust	H-Trust	PeerTrust	WTR	BubbleTrust
Camouflage	0.17	0.19	0.09	N/A	0.16
Full Collusion	N/A	8.58	2.78	2.72	20.31
Evaluator Collusion	N/A	9.67	9.06	9.12	36.47
Spies	1.95	5.79	2.23	2.13	29.87
Evaluator Spies	5.87	5.74	5.72	5.69	36.23
Malicious Spies	N/A	5.65	5.68	5.56	31.63

Table 6. MaliciousBenefit for different malicious strategies and TMSs.

Strategy	EigenTrust	H-Trust	PeerTrust	WTR	BubbleTrust
Camouflage	0.17	0.19	0.09	N/A	0.16
Evaluator Collusion	N/A	6.79	6.36	6.41	25.64
Spies	0.10	0.24	0.11	0.09	1.96
Evaluator Spies	2.85	2.73	2.75	2.74	17.73
Malicious Spies	N/A	2.67	2.68	2.63	14.95

The attacker most likely uses a strategy which has the best price/performance ratio. For instance, in the PeerTrust the most successful strategy is Evaluator Collusion but it is very expensive (above 9), better choice is Full Collusion with success ratio 0.99 and cost only 2.78. The Camouflage strategy is relatively efficient; although it has low a success ratio in the most TMSs, it is compensated by its very low price. In the BubbleTrust, all strategies have cost above 20 (except Camouflage) and the most expensive strategy (Evaluator Collusion) has almost 37. This is significantly higher value than the other TMSs have.

Table 6 shows MaliciousBenefit of malicious strategies which have some beneficial transactions. Again, MaliciousBenefit of the strategies with no measurable MSR is infinite and the cells contain ‘N/A’.

The strategies like Evaluator Collusion, Evaluator Spies and Malicious Spies have always more beneficial transactions than bogus ones. Strictly speaking, the designation of the collective as malicious is no longer suitable. The attackers, whose primary goal is to destroy the network functionality for other peers, probably would not choose malicious strategy with a high MaliciousBenefit. But attackers desired to spread their malicious services at any cost do not bother with MaliciousBenefit.

3.3 Influence of Simulation Settings

We have tried different simulation settings. We have adjusted the number of nodes in the network while preserving the ratio of malicious nodes. We have made the following observation: increasing the number of nodes does not affect the MaliciousSuccessRatio. The reason is that each TMS can handle only a limited number of nodes in the calculation of ratings. A similar limitation can be found in all TMSs. The information from nodes which are very distant in a trust chain is negligible. On the other hand, the results change if we decrease the number of nodes. This change can be in both

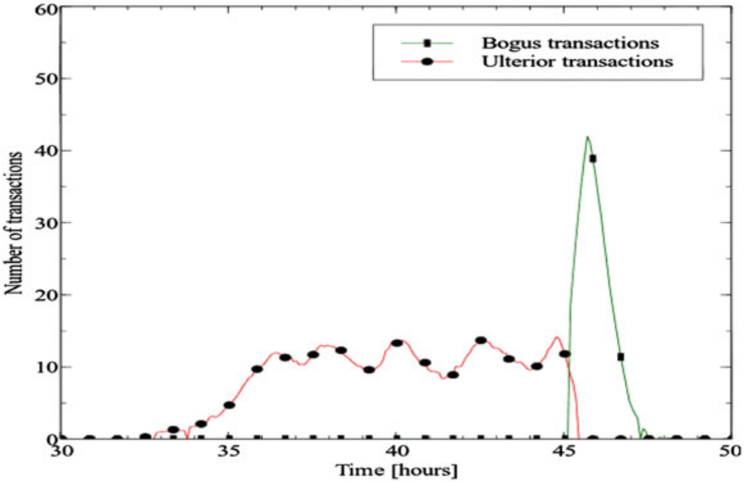

Fig. 1. Rehabilitation after treason in BubbleTrust.

directions dependent on the TMS and the malicious strategy. In this case the TMS has to rely on information from a smaller number of nodes than it expects (Fig. 1).

Next, we have altered the ratio of malicious nodes. Figure 2 shows the results for BubbleTrust. As we can see, the malicious success increases with the ratio of malicious nodes in the network. BubbleTrust resists relatively well even in the network with more than 50 % of malicious nodes. In our tests we stayed at 40 % because it is very unlikely that the overlay network beneath the P2P application can handle the situation in which half of the peers are malicious. The defence techniques described in

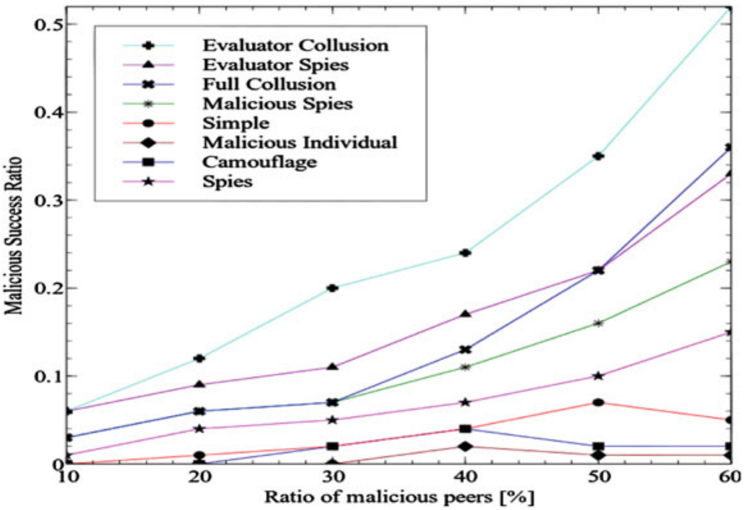

Fig. 2. Ratio of malicious peers on Malicious Success Ratio in BubbleTrust.

2.1 assume that only a small fraction of nodes is malicious. In fact, 40 % already causes big problems.

3.4 Result Summary

H-index calculation used in H-Trust proved to be vulnerable to traitors. It takes too long to detect traitors and malicious peers are rehabilitated too quickly. The system WTR permits the highest number of bogus transactions from all tested TMSs, but it is followed closely by PeerTrust and HTrust. EigenTrust has better results than H-Trust, WTR and PeerTrust but it has advantage in the form of pre-trusted peers.

Our tests proved that it is very difficult to resist against the sophisticated malicious techniques. Especially the calculation of the evaluator rating is susceptible to rigging. The previously published TMSs do not pay as much attention to the evaluator rating as they pay to the provider rating. This must be changed if the TMS should be resistant against the Evaluator Collusion or the Evaluator Spies.

The best TMS in our comparison is BubbleTrust. It has the shortest treason detection time, the longest rehabilitation time and allows only 28 % of bogus transaction under the most successful malicious strategy. As far as we know, it is the only one TMS using global experience as feedback verification.

4 Conclusion

In this paper, using simulation framework called P2PTrustSim we compared trust management systems against different malicious strategies. We also proposed several efficiency criteria which can be evaluated using this framework. We analysed known malicious tactics and suggested three new collective malicious strategies against the most representative systems for each type of TMS. We can expect that malicious peers working in a collective will try to use the most effective strategy against TMS currently used in the network. Therefore, the quality of TMSs has to be assessed according to the most successful malicious strategy. Nevertheless, other properties have to be taken into account too; e.g. the cost connected with the malicious strategy can exceed the potential benefit for malicious peers. The results indicate that only the BubbleTrust is resistant against all considered malicious strategies; it is, therefore, the best choice for deployment in the secured P2P networks.

Acknowledgment. This work was supported in part by grants 204/13/08195 and SVV-2013-267312.

References

1. Kamvar, S.D., Schlosser, M.T., Garcia-Molina, H.: The Eigentrust algorithm for reputation management in P2P networks. In: WWW'03: Proceedings of the 12th International Conference on World Wide Web, pp. 640–651. ACM Press (2003)
2. Xiong, L., Ling, L.: PeerTrust: supporting reputation-based trust for peer-to-peer electronic communities. IEEE Trans. Knowl. Data Eng. **16**, 843–857 (2004)

3. Huanyu, Z., Xiaolin, L.: H-Trust: a group trust management system for peer-to-peer desktop grid. J. Comput. Sci. Technol. **24**, 447–462 (2009)
4. Bonnaire, X., Rosas, E.: WTR: a reputation metric for distributed hash tables based on a risk and credibility factor. J. Comput. Sci. Technol. **24**, 844–854 (2009)
5. Novotny, M., Zavoral, F.: BubbleTrust: a reliable trust management for large P2P networks. In: Meghanathan, N., Boumerdassi, S., Chaki, N., Nagamalai, D. (eds.) CNSA 2010. CCIS, vol. 89, pp. 359–373. Springer, Heidelberg (2010)
6. Novotny, M., Zavoral, F.: Resistance against malicious collectives in BubbleTrust. In: The 12th International Conference on Parallel and Distributed Computing, Gwangju, Korea (2011)
7. Hoffman, K., Zage, D., Nita-Rotaru, C.: A survey of attack and defense techniques for reputation systems. ACM Comput. Surv. **42**(1), 1–31 (2009)
8. Marti, S., Garcia-Molina, H.: Taxonomy of trust: categorizing P2P reputation systems. Comput. Netw. **50**, 472–484 (2006)
9. Selvaraj, C., Anand, S.: Peer profile based trust model for P2P systems using genetic algorithm. Peer-to-Peer Netw. Appl. **4**, 1–12 (2011)
10. Suryanarayana, G., Taylor, R.N.: A survey of trust management and resource discovery technologies in peer-to-peer applications. Technical report, UC Irvine (2004)
11. Bonnaire, X., Rosas, E.: A critical analysis of latest advances in building trusted P2P networks using reputation systems. In: Weske, M., Hacid, M.-S., Godart, C. (eds.) WISE Workshops 2007. LNCS, vol. 4832, pp. 130–141. Springer, Heidelberg (2007)
12. Druschel, P., Rowstron, A.: PAST: a large-scale, persistent peer-to-peer storage utility. In: Proceedings of the Eighth Workshop Hot Topics in Operating Systems (2001)