

Trust-Based Taxonomy for Free Riders in Distributed Multimedia Systems

Farag Azzedin

Information and Computer Science Department
King Fahd University of Petroleum & Minerals, Saudi Arabia
fazzedin@kfupm.edu.sa

ABSTRACT

The ever growing explosion in technological advancements is paving the way to the expansion of multimedia applications. To fully utilize such applications, some issues related to decentralization must be addressed. In this paper, we focus on the problem of free riding in decentralized collaborative environments. We propose a novel taxonomy of free riders in multimedia systems based on trustworthiness. We believe that trustworthiness is a vital dimension that should be considered when identifying free riders. We also propose a new mechanism to filter out and isolate free riders. Our extensive simulation experiments show that our proposed algorithm is reasonably successful in identifying free riders in multimedia-based systems.

KEYWORDS: multimedia systems, free riders, peer-to-peer, distributed computing.

1. INTRODUCTION

Peer-to-peer (P2P) architectures have proven [1] the feasibility of global multimedia content. BitTorrent [8], the large scale standard P2P technology for the delivery of rich media on the Internet, has recently joined forces with Orb Networks [2], the streaming media leader, to bundle BitTorrent software client applications and the Orb PC applications together. This joint effort will make BitTorrent downloads feasible and enjoyable on any connected device. Multimedia streaming such as PPStream [7], SOPCast [12], PPLive [11], as well as live TV over P2P network (P2PIPTV) are emerging as successful applications of the growing explosion of the digital multimedia-based era. This era is moving towards a large scale multimedia content sharing.

The large scale multimedia content is a collaborative effort from users willing to share and contribute to achieve collectively and successfully a joint goal. Unfortunately, if users start to consume resources more than they contribute, then the burden of uploading will be on the altruistic users. These altruistic users might be very few and this will create bottlenecks resulting in degrading the *quality-of-service* (QoS). In both live streaming and *video-on-demand* (VoD), a minimal download speed is needed to sustain playback. As such, if users do not contribute, distributing multimedia content over the Internet may not survive. Also, burdening the collaborative users with huge requests might force them to be non-collaborative resulting in adding insult to injury. Furthermore, non-collaborative users negate the advantages of distribution and hence convert a distributed architecture to few scattered centralized nodes.

Non-collaborative users, who do not contribute to the community, are referred to as *free riders* in the literature. The phenomenon of free riding does not just create a tedious task that can be bothersome for the contributors, but also can create a potential performance bottleneck. As a consequence, responses to queries can experience a longer delay. In a collaborative file sharing system, 66% of users do not share files (i.e., free riders) and the top 25% users account for 99% of all downloads [10]. Free riding must be tackled before P2P-based technology can be used to provide the necessary substrate to multimedia applications and services so that they can be redeployed over P2P networks.

1.1. Motivation

We were motivated by the fact that existing free riders' detection algorithms classify a node as a free rider based on the node's contribution versus consumption as shown in Equation 1. Equation 1 will consider a node that contributes untrustworthy content or resources as not a free rider. In this

paper, we argue that an untrustworthy node (i.e., a node intentionally providing corrupted resources) fits the classification of a free rider.

$$\Omega = \frac{\text{consumption}}{\text{contribution}} \quad (1)$$

The consumption-to-contribution ratio (i.e., Ω) is a reasonable measure to detect a free rider under the assumption that the node is trustworthy. A node is trustworthy if it does not contribute content or resources that inflict damage to other nodes. The intent of an untrustworthy node is to inject corrupted resources and hide its malicious behavior. To pursue this vicious intent, an untrustworthy node can exploit Ω . This can be done by injecting the corrupted resources such that the malicious node is considered to be contributing and hence its Ω is satisfactory.

In fact, when an untrustworthy node provides corrupted resources, it is actually consuming resources from other nodes without contributing anything. If node x is providing a corrupted file and node y is interested in downloading that file, then node y will allocate some of its resources to download the corrupted file. These wasted allocated resources as well as the time node y spent to download the corrupted file are all indirectly consumed by node x . How does y detect that the file is corrupted, is outside the scope of this paper. Detection can be done online or offline using post mortem analysis techniques.

In addition, node x 's behavior negates the advantages of distribution and sharing. Node x also contributes negatively to the active role of contributing nodes as they become overloaded and this might encourage selfishness. An untrustworthy contributor fits the classification of free riders because such a node is not contributing to the community, such a node is allowed to consume since its Ω is satisfactory, and such a node discourages trustworthy contribution. Therefore, an untrustworthy node does not only fit the classification of a free rider but also introduces new challenges to the success of multimedia services and applications.

1.2. Contributions

We propose a new mechanism to filter out and prevent free riders from degrading and hence bringing collaborative environments such as multimedia systems to its knees. In our mechanism, we use the notion of trustworthiness. Trustworthiness is a notion that describes the behavior of a node and it is a vital component in any Internet-based transaction [3, 9]. Throughout this paper, we refer to the node that

wants to consume resources as a *source node* whereas the node that provides resources as a *target node*.

The contributions of this paper is three folds: (a) draw the attention and raise the awareness to the fact that trustworthiness is an important attribute for identifying free riders, (b) provide a trust-based taxonomy for free rider, and (c) propose and evaluate a novel activity-based algorithm to identify free riders in multimedia-based systems. To the best of our knowledge, no existing literature tackles this issue, which we believe is a vital dimension that should be considered when identifying free riders.

2. FREE RIDERS TAXONOMY

Free riders attracted the attention of many research groups. Current research done on streaming systems [3, 4, 7] as well as research on P2P and VoD [1, 2] have used a contribution-based taxonomy to classify free riders. These researchers have assumed that free riders are trustworthy and they are either selfish, consumers, or droppers. A selfish free rider is a node that does not share its own resources, does not replicate resources received from others, but it routes messages and generates normal traffic. A consumer free rider is a node that might share its own resources, does not replicate resources received from others, but it routes messages and generates high traffic. A dropper free rider is a node that does not share its own resources, does not replicate resources received from others, generates normal traffic, but does not route messages.

All these researchers use the ratio consumption over contribution as a performance metric to identify a free rider, where consumption might be number of files downloaded or query hit messages and contribution might be number of uploads or query messages sent to neighbors. We argue that a free rider might be untrustworthy (i.e., malicious) and hence contribute its malicious content or even contribute modified versions of replicas received from other nodes. An untrustworthy node will not be identified as a free rider using the ratio performance metric. As such, an untrustworthy node will be given the chance to damage a collaboration system and will not be labeled as a free rider.

We propose a taxonomy of free riders based on trustworthiness as shown in Figure 1. A free rider can be selfish meaning that the node is not contributing at all but only using or consuming resources. On the other hand, a free rider can contribute. Taking trustworthiness into consideration, a contributor free rider can be either trustworthy or untrustworthy. A trustworthy or untrustworthy contributor free rider can easily detect if it is identified because its downloads, for

example, are becoming slow since it is given smaller bandwidth than it used to get previously.

A contributor free rider can adapt a flip-flop phenomenon. Once it detects that it was identified as a free rider, it will start to contribute and then it will not contribute until it is again identified as a free rider. This cycle will repeat and the free rider will conditionally contribute and not for the sake of contribution. An untrustworthy free rider will negatively contribute to the collaborative system and will not be identified as a free rider. It can inject its malicious content based on conditions such as a time frame (i.e., temporal), when needed (i.e., on need), infecting just a specific regions or group of users (i.e., spacial), or worse infecting everyone any time (i.e., always).

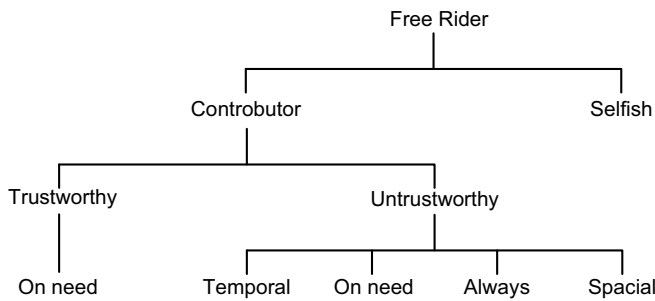


Figure 1: Trust-based taxonomy for free riders.

3. LITERATURE REVIEW

With the recent proliferation of multimedia applications and the deployment of ad-hoc networks, it became possible for peers to communicate and share resources using video collaboration and e-meeting services from almost any location at anytime [6]. In multimedia applications, trust is recognized by the research community as an important factor that contributes to the success such applications.

A contribute-to-consume treatment is proposed in [1] for a P2P VoD system, where a free rider is identified based on its ratio of consumption over contribution. Again, the authors assumed that free riders are trustworthy and therefore well-behaved. The proposed algorithm is inspired by BitTorrent [8] and the performance evaluation shows that the algorithm provides good performance as long as the all peers are trustworthy and the shared videos are short, which the authors claim that short videos dominate the current VoD traffic on the Internet.

Incentives for free riders in streaming live multimedia systems are proposed in [5]. The authors' aim is to deploy live

multimedia streaming over P2P systems and hence uncovering distributed approaches to manage multimedia-based applications. In their approach, the authors emphasize on the temporal importance that exists in live multimedia streams. Therefore, block a in a stream, needs to be consumed before block $a + 1$ and hence receiving block a is more important than receiving block $a + 1$. This temporal property might make a certain node not transmitting, not because it is non-collaborative, but because the temporal constraints makes the transmission pointless.

P2P systems have proven its success due to their attractive properties such as self-organizing, scalability, distributed resource sharing, self-autonomy, and more [13, 14]. Unfortunately, issues such as free riding is an important concern pertaining to resource sharing. Many researchers proposed solutions, to detect and discourage free riding, ranging from game theory to providing incentives [13, 14].

The authors in [13] used an algorithm to track free riders in file sharing applications. The authors used the popular choking algorithm [8] and how free riders affect the file sharing network. The authors excluded the concept of malicious nodes (i.e., they assumed that all nodes are trustworthy) and they used the ratio of download over upload to detect and isolate free riders.

In [14], the authors addressed user incentives in a P2P storage system. In their paper, the authors proposed that each node will use resources based on its contribution. Using game theoretical analytical model to describe free riders reactions, the authors devised a scheme to over incentives for cooperation and model users' behavior in relation to resource availability.

4. THE PROPOSED APPROACH

The aim of the proposed free riders' filtering algorithm is to identify and isolate free riders. In our algorithm, every peer maintains a black list. Each time a source peer interacts with a trustworthy target peer, two things can happen. If the target peer is trustworthy, the source peer globally posts an appraisal. This global posting can be easily established in structured P2P infrastructures such as P-Grid [15] and the information is kept by a third-party peer, e.g., the storage peer. On the other hand, if the target peer is untrustworthy, the source peer will not post anything into the global list but it will mark the untrustworthy peer in its black list. In other words, a source peer can appraise but it can not complain. Using its black list, a source peer will not interact with a black listed target peer.

The algorithm uses an activity-set model to monitor the ac-

tiveness of each peer. By activeness, we mean the trustworthy contribution of a peer. Let Δ be the activity-set window. The idea is to examine the number of times peer a got appraised in the current activity-set window Δ_c^a and the previous activity-set window Δ_p^a and hence the trustworthy contribution of peer a (TC_a) can be measured according to equation 2.

$$TC_a = \begin{cases} 1 & \text{if } \Delta_c^a > AC \times \Delta_p^a \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

In equation 2, AC is the activity constant, where $0 \leq AC \leq 1$ and is used to tune the comparison operation between Δ_c^a and Δ_p^a . If the TC_a is 0, then peer a is considered to be a free rider. Otherwise, it is not a free rider.

Equation 2 will consider peer a with constant trustworthy contribution x as not a free rider even if $x = 1$. That is, a contributes only once in each Δ . To solve this problem, our algorithm uses a contribution factor of peer a (CF_a) that is calculated as follows:

$$CF_a = \frac{x}{N_{num}} \quad (3)$$

As stated in [9], when peers download and exchange conventional multimedia content, e.g., MP3 files for popular music, the content will likely be uniformly distributed. That is, if Δ is determined based on the number of transactions (T_{num}), then the average number of transactions (N_{num}) can be easily calculated based on Equation 4, where P_{num} is the number of peers in the environment.

$$N_{num} = \frac{T_{num}}{P_{num}} \quad (4)$$

Now, if CF_a is greater than a threshold, then the contribution factor of peer a is acceptable and we can proceed to Equation 2. Otherwise, peer a is considered a free rider.

For further illustration of how the proposed free riders' filtering algorithm works, Algorithm 1 shows how the appraisals are posted and Algorithm 2 shows how these appraisals are utilized to filter free riders.

Algorithm 1 shows the appraisals posting process performed by peer a . Line 2 initializes peer a 's black list. In line 4, peer a chooses a target peer t . Lines 5 and 6 state that if t is not black listed by a , then a will perform the transaction with t . Lines 7 to 12 states that if t is not in the black list of peer a , then a will interact with t . If a finds that t is trustworthy (i.e., shares trustworthy content), then a will

Algorithm 1 Appraisals posting algorithm. This algorithm is run by peer a

```

1: procedure ACTIVITYBASED
2:    $BL_a \leftarrow \{\}$ 
3:   repeat
4:     select a target peer  $t$ 
5:     if  $t \notin BL_a$  then
6:       perform the transaction with  $t$ 
7:       if  $t$  is trustworthy then
8:         peer  $a$  posts an appraisal for  $t$ 
9:       else
10:         $BL_a \leftarrow BL_a + t$ 
11:      end if
12:    end if
13:  until peer  $a$  is not interested in performing transactions
14: end procedure

```

Algorithm 2 The filtering algorithm. This algorithm is run by the storage peer of peer a

```

1: procedure THE FILTERING ALGORITHM
2:    $x = \Delta_c^a$ 
3:    $y = \Delta_p^a$ 
4:    $CF_a = \frac{x}{N_{num}}$ 
5:   if  $CF_a$  is satisfactory then
6:     Compute  $TC_a$  based on  $x$  and  $y$ 
7:     if  $TC_a == 1$  then
8:       peer  $a$  is not a free rider in  $\Delta_c$ 
9:     else
10:      peer  $a$  is a free rider in  $\Delta_c$ 
11:    end if
12:  end if
13: end procedure

```

post an appraisal with the storage peer of t . Otherwise, a will black list t .

For further illustration, the steps of the filtering algorithm are shown in Algorithm 2. This algorithm can be run by each storage peer in the P-Grid [15] infrastructure. In other words, if peer y is the storage peer for peer a , then y will determine whether a is a free rider or not. Algorithm 2 will start by computing the number of appraisals posted for a specific Δ . In lines 2 and 3, the storage peer gets the number of appraisals for the current and the previous activity-set windows. Then in lines 4 to 10, the contribution factor is calculated. If the contribution factor is satisfactory, peer a is identified whether it is a free rider by calculating TC_a as shown in equation 2.

5. PERFORMANCE EVALUATION

5.1. Performance Metrics

We use the following performance metrics to evaluate the effectiveness of the free riders' filtering algorithm:

1. **Detection Tolerance.** The free riders' filtering process ideally distinguishes free riders from positive contributors so that positive contributors can be recognized and rewarded for their effort in boosting and making multimedia applications a success. For the detection tolerance, we measure the following: (a) The percentage of free riders detected as free riders (true positives, i.e., T+) and (b) The percentage of non-free riders mis-detected as free riders (false positives, i.e., F+).
2. **Free riding Impact.** Free riders would affect how the source peer makes its decision. We measure the sensitivity of free riding towards the decision making by computing: (a) The percentage of lost trustworthy transactions. Lost trustworthy transactions are due to F+ and (b) The percentage of committed untrustworthy transactions. Committed untrustworthy transactions are due to not identifying untrustworthy contributors.

5.2. Simulation Setup

We use 2048 peers and the simulation runs until in average each peer performs 40 transactions. The source and the target peers for each transaction are randomly generated from $[1, N]$, where $N = 2048$. The transaction will be conducted if the target peer is predicted as a positive contributor. If the transaction is committed, the source peer globally will provide an appraisal based on the trustworthiness of the target peer. In the simulation, we used P-Grid [15] to post the appraisals. If the target peer is untrustworthy free rider, the target peer will be marked in the source peer's black list.

We experiment with three types of free riders, namely selfish (S), trustworthy contributors (TC), and untrustworthy temporal contributors (UTC). For each type, we use different proportions of free riders, i.e., 20%, 40%, 60%, and 80%. All results shown are an average of 10 simulation runs. The size of 2048 nodes provides a confidence interval of 3.05% and 4.01% at confidence level of 95% and 99%, respectively, if the population size is 118,925 nodes (the size of Gnutella network in February 2006 [16]). If the population size is assumed to be 1 million nodes, the confidence interval changes slightly to 3.07% and 4.03% at confidence

level of 95% and 99%, respectively. As for the average number of transactions performed by peers, it indicates the length of the simulation. A free riders' filtering mechanism should be able to detect free riders quickly so that they do not have many opportunities to pollute the environment. The activity constant AC is set to 0.66.

The simulation starts with no transactions performed so there are no appraisals available. We use optimistic stranger policy i.e., we assume that new peers are not free riders. If a trustworthy source peer transacts with an untrustworthy contributor, it will detect the malicious content and will not offer it to the community. For the simulation, we assume that all peers are honest. That is, a source peer will not appraise an untrustworthy contributor and will appraise a trustworthy contributor.

6. RESULTS AND DISCUSSION

6.1. Detection Tolerance

With 20% free riders as shown in Figure 2(a), the detection rate is between 96% and 100% for S , TC , and UTC free riders. The detection rate falls a little for the TC 's because the peer's transactions may fall on either side of Δ boundary (i.e., a TC free rider a gives an appraisal in time slot x but receives an appraisal in time slot $x + 1$. In that case it will not be detected as a free rider. In Figure 2(b), the mis-detection rate for the three types of free riders is very low and ranges from 1% to 2%. From these results, we see that our algorithm maximizes T+ and minimizes F+.

Figure 3(a) shows the detection rate for 40% free riders. The detection rate is 100% for S free riders since this type of free riders uses resources and does not share resources. So, the S free riders are always source peers and never target peers. Since an appraisal is made by the source peer for the target peer after the transaction is successful, this type of free riders will never get an appraisal and hence is easily detected. For the UTC free riders, the detection rate is also 100%. Since the UTC free riders contribute untrustworthy content, they will never get an appraisal. As such, this type of free riders is detected by our algorithm as well. On the other hand, the detection rate for the TC free riders ranges from 96.70% to 98.53% because the TC basically adapts a flip-flop phenomenon based on need. In spite of that, the detection rate of our algorithm of this type of free riders is more than 96%.

The misdetection rate when free riders are 40% is shown in Figure 3(b). The misdetection rate is very low and ranges from 0.1% to 0.7% because the non-free riders are only 60%

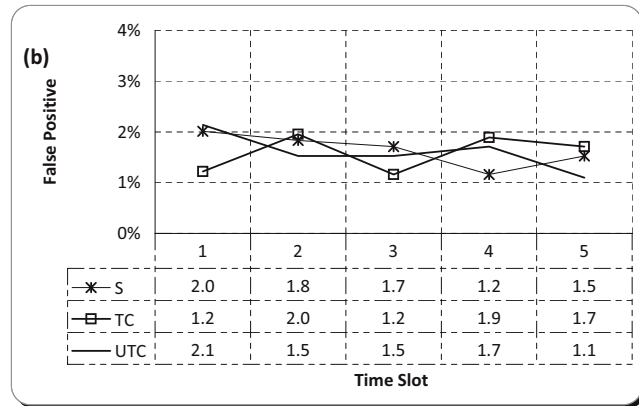
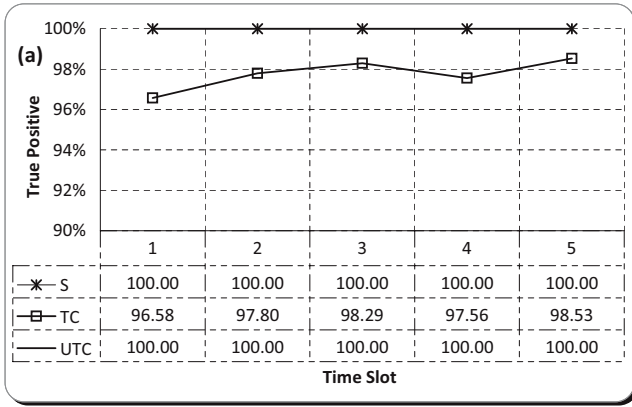


Figure 2: Detection tolerance using 20% free riders: (a) Free riders detected as free riders and (b) Non-free riders misdected as free riders.

compared to 80% in Figure 2(b). As it shows in Figures 2(b), 3(b), 4(b), 5(b) there is a trend in the misdetection rate. As the number of non-free riders gets smaller, the misdetection rate also gets smaller. These results validate our simulator and logically follow the definition of the misdetection rate which is the percentage of non-free riders misdected as free riders.

The detection rate for 60% and 80% free riders are shown in Figure 4(a) and 5(a), respectively. For both Figures, the detection rate for *S* as well as *UTC* is 100%. This for the same reason explained for Figures 2(a) and 3(a). On the other hand, detecting the *TC* free riders is more difficult since they change their behavior in a flip flop manner. This becomes apparent in the case of 80% free riders where we notice that their detection rate falls to about 91.03%.

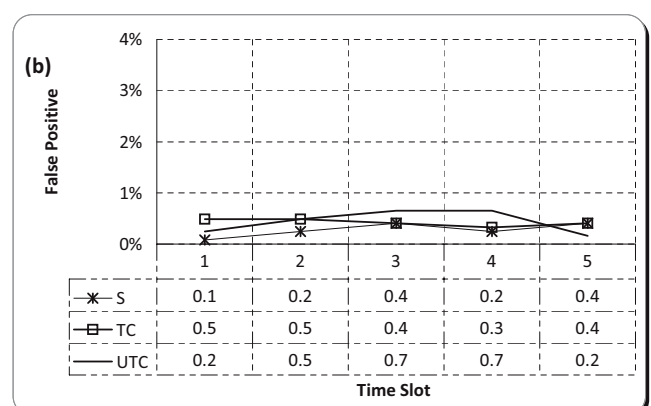
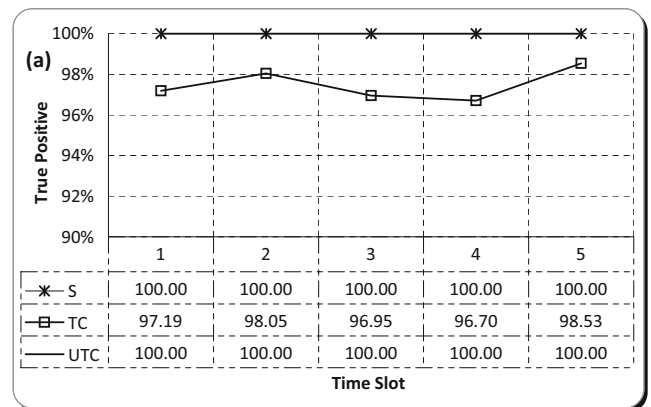


Figure 3: Detection tolerance using 40% free riders: (a) Free riders detected as free riders and (b) Non-free riders misdected as free riders.

6.2. Free Riders Impact

In this set of experiments, we examine the effect of free riders on the decision making process made by a source peer. We focus on the *UTC* free riders because their impact is the worst since this type of free riders are injecting untrustworthy content and trying to pollute the network. As such, they infect the environment with malicious content for nefarious motives.

In Figure 6(b), the committed untrustworthy transactions are directly related to the detection rate shown in the previous Figures. As shown in Figures 2(a), 3(a), 4(a), the *T+* for *UTC* is 100%. This means that all free riders are detected and hence no untrustworthy transaction is committed. This fact is confirmed in Figure 6(b).

In Figure 6(a), the lost trustworthy transactions are also directly related to the misdetection rate. The committed un-

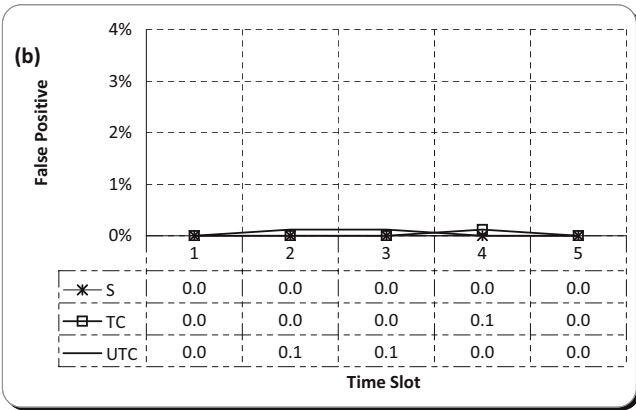
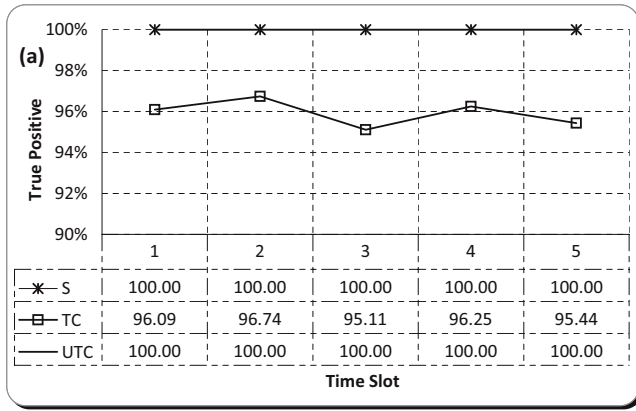


Figure 4: Detection tolerance using 60% free riders: (a) Free riders detected as free riders and (b) Non-free riders misdetected as free riders.

trustworthy transactions are 0 as shown in Figure 6(b) because the untrustworthy transactions will only happen if a source peer transacts with a *UTC* free rider. A source peer will only transact with peers having appraisals greater than 0.66 of the previous Δ . The untrustworthy contributors' appraisals are never going to increase and no source peer will commit a transaction with them.

7. CONCLUSION AND FUTURE WORK

To deploy multimedia applications on distributed architectures such as P2P systems, many issues need to be resolved. In this paper, we investigated the free riding problem. Currently, free riding is tackled based on tit-for-tat strategies. With such strategies, an untrustworthy free rider, who consumes but injects malicious content, is undetected. This leads to the corruption and degradation of the distributed multimedia applications because peers will refrain from us-

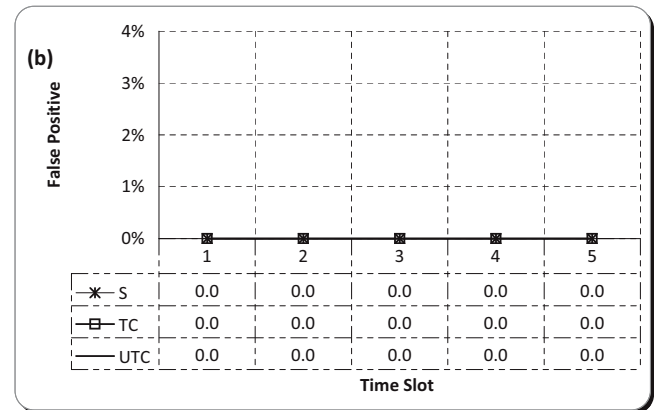
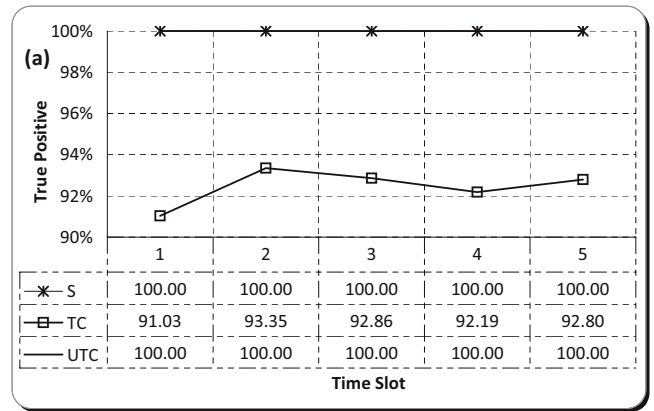


Figure 5: Detection tolerance using 80% free riders: (a) Free riders detected as free riders and (b) Non-free riders misdetected as free riders.

ing the environment due to the lack of trust.

We argue that trustworthiness is a crucial attribute for free riders since an untrustworthy free rider can maintain its contribution factor by inflicting damage to the P2P community with its malicious content. Therefore, such untrustworthy contributors should be considered as free riders because they do not contribute but actually pollute and discourage contribution.

We proposed a novel free riders' filtering algorithm that detects selfish, trustworthy and untrustworthy free riders. Our simulation results show that the proposed algorithm identifies most of free riders even if the free riders are 80% of the whole population. The results also show that our algorithm reduces the impact of free riders by minimizing the lost trustworthy transactions and the committed untrustworthy transactions. For future work, we are working on injecting dishonest free riders that lie. By considering honesty,

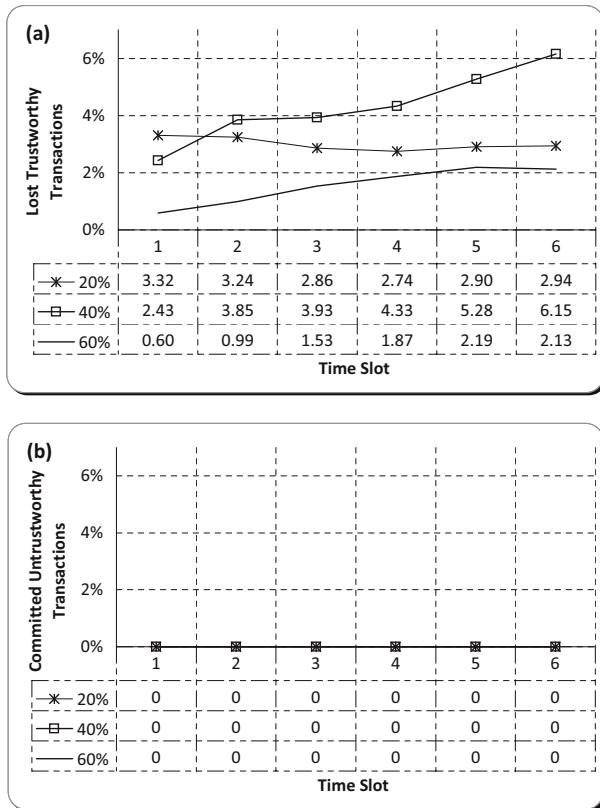


Figure 6: Free riders impact: (a) Lost trustworthy transactions and (b) Committed untrustworthy transactions.

we can evaluate our algorithm’s resilience to free riders bad-mouthing or colluding.

ACKNOWLEDGEMENTS

The author would like to thank all anonymous reviewers for their helpful comments , resulting in a significant improvement in the quality of this manuscript. This research is supported by King Fahd University of Petroleum and Minerals under Research Grants IN080430 and FT060028.

REFERENCES

[1] Mol, J., Pouwelse, J., Memulpolder, M., Epema, D., Sips, H.: “Give-to-Get: free-riding-resilient video-on-demand in P2P systems”; Proc. Multimedia Computing and Networking 2008, San Jose (Jan 2008).

[2] Orb Networks, Inc. <http://www.orb.com/>.

[3] Azzedin, F., Ridha, A: “Feedback Behavior and Its Role in Trust Assessment for Peer-to-Peer Systems”; Telecommunication Systems. 2009. DOI: 10.1007/s11235-009-9263-9.

[4] Chatzidrossos, I., Fodor, V.: “On the effect of free-riders in P2P streaming systems”; Telecommunication Networking Workshop on QoS in Multiservice IP Networks,(Feb 2008).

[5] Silverston, T., Fourmaux, O., Crowcroft, J.: “Towards an incentive mechanism for peer-to-peer multimedia live streaming systems”; the Eighth International Conference on Peer-to-Peer Computing, (Feb 2008).

[6] Chatzidrossos, I., Fodor, V.: “On Technologies for Ubiquitous Multimedia Communication”; PhD Thesis, Lulea University of Technology, Sweden, (2006).

[7] ppstream TV. <http://www.ppstream.com/>.

[8] Shah, P., Paris, F.: “Peer-to-Peer Multimedia Streaming Using BitTorrent”; the IEEE Performance, Computing, and Communications Conference, (April 2007) 340-347.

[9] Mihaela S., Deepak S. T., Ritesh S.: “Stochastic Optimization for Content Sharing in P2P Systems”; IEEE TRANSACTIONS ON MULTIMEDIA 1, 1 (Jan 2008).

[10] Tang, Y., Wang, H., Dou, W.: “Trust based incentive in P2P network”; the IEEE International Conference on E-Commerce Technology for Dynamic E-Business, (2004) 302-305.

[11] PPLive. <http://www.pplive.com/>.

[12] SopCast. <http://www.sopcast.com/>.

[13] Obele, B., Ukaegbu, A., Kang, M.: “On Tracking Free-Riders in P2P Networks”; Proc. the 11th International Conference on Advanced Communication Technology, 2009.

[14] Maille, P., Toka, L.: “Managing a Peer-to-Peer Data Storage System in a Selfish Society”; IEEE Journal on In Selected Areas in Communications 26, 7 2008, 1295-1301.

[15] Aberer, K., Datta, A., Hauswirth, M.: “P-Grid: Dynamics of self-organization processes in structured P2P systems”; Proc. Peer-to-Peer Systems and Applications, Lect. Notes in Comp. Sci. 3845, Springer, Berlin. 2005.

[16] Xie, C., Chen, G., Vandenberg, A., Pan, Y.: “Analysis of Hybrid P2P Overlay Network Topology”; Computer Communications 31, 2 (Feb. 2008), 190-200.