Computer Networks 57 (2013) 1019–1036

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

SimplyRep: A simple and effective reputation system to fight pollution in P2P live streaming

Alex Borges Vieira [a],*, Rafael Barra de Almeida [a], Jussara Marques de Almeida [b], Sérgio Vale Aguiar Campos [b]

[a] Computer Science Department, Universidade Federal de Juiz de Fora, Brazil
[b] Computer Science Department, Universidade Federal de Minas Gerais, Brazil

ARTICLE INFO

Article history:
Received 28 April 2012
Received in revised form 30 November 2012
Accepted 1 December 2012
Available online 8 December 2012

Keywords:
P2P
Live streaming
Attack
Reputation

ABSTRACT

Peer-to-Peer (P2P) streaming has become a popular platform for transmitting live content. However, due to their increasing popularity, P2P live streaming systems may be the target of user opportunistic actions and malicious attacks, which may greatly reduce streaming rate or even stop it completely. In this article, we focus on a specific type of attack called content pollution, in which malicious peers tamper or forge media data, introducing fake content before uploading it to their partners in the overlay network. Specifically, we present a new decentralized reputation system, named SimplyRep, that quickly identifies and penalizes content polluters, while incurring in low overhead in terms of bandwidth consumption. We evaluate our method with both simulation and experiments in PlanetLab, comparing it against two previously proposed approaches, namely, a centralized black list and a distributed reputation system, in various scenarios. Our results indicate that SimplyRep greatly outperforms the two alternatives considered. In particular, both black list and the distributed reputation method perform poorly when polluters act jointly in a collusion attack, reaching a data retransmission overhead (triggered by polluted chunks received) of 70% and 30%, respectively, whereas the overhead experienced by SimplyRep is at most 2%. Our results also show that SimplyRep is able to quickly isolate almost all polluters under a dissimulation attack, being also somewhat robust to a whitewashing attack, although the latter remains a challenge to effective P2P streaming.

1. Introduction

P2P live streaming is becoming increasingly popular. Indeed, some important TV channels already broadcast their live content on the Internet using P2P technology. For example, CNN relied on a P2P platform to assist CDNs on the live transmission of the Barack Obama's inauguration speech, which is seen as one of the largest live video event in the history of the Internet. In fact, the live transmission exceeded 1.3 million simultaneous streams, more than half of which were delivered over a P2P network.[1]

Despite the recent success, P2P live streaming protocols are very susceptible to malicious attacks, which can hinder their adoption as an alternative architecture to traditional client server protocols. Moreover, security issues in P2P live systems are more challenging than in other P2P applications because live transmissions are more vulnerable to QoS fluctuations due to their strict delay constraints.

* Corresponding author. Tel.: +55 32 2102 3311.
E-mail addresses: alex.borges@ufjf.edu.br (A.B. Vieira), rafael.barra@ice.ufjf.edu.br (R.B. de Almeida), jussara@dcc.ufmg.br (J.M. de Almeida), scampos@dcc.ufmg.br (S.V.A. Campos).

[1] This statistic was extracted from an article entitled "CNN: Inauguration P2P Stream a Success, Despite Backslash", by Janko Roettgers, published at the GigaOM portal on February 7th 2009, and available, as of August 2012, at http://gigaom.com/video/cnn-inauguration-p2p-stream-a-success-despite-backlash/.

1389-1286/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2012.12.001

One particular malicious attack that threatens P2P streaming is content pollution, where malicious peers (polluters) tamper or forge media data, introducing fake content and/or advertisements, before uploading it to their partners. Moreover, non-malicious (i.e., legitimate) peers cannot easily distinguish between polluted and genuine content before watching it. Thus, unaware of the legitimacy of a piece of content they received, these peers end up forwarding polluted content to their own partners, acting as passive polluters. Without proper defense mechanisms, polluted data spreads quickly over the P2P network, causing significant impact on the system in terms of resource (particularly bandwidth) consumption and streaming quality [1–3].

Various strategies to fight pollution attacks have been proposed in the literature, including black listing [1,4], hash-based signatures [1,5,6,3], data encryption [1,3] and reputation systems [2,7,8]. In spite of that, most popular applications do not use any protocol or data encryption strategy [1], possibly because existing techniques may significantly increase both processing and communication overhead (and thus peer resource requirements), as well as media startup latency. Thus, most current P2P live applications remain vulnerable to pollution attacks.

In this article, we tackled the problem of fighting content pollution in P2P live streaming systems by proposing a new simple and decentralized reputation system. Our system, called SimplyRep, aims at quickly detecting peers that upload polluted content, here referred to as content polluters. Unlike previous decentralized reputation systems [2,9–15], in SimplyRep, a peer reputes its partners based solely on the rate of polluted/damaged data it received from them. To that end, SimplyRep relies on any existing method to detect polluted content once it is received [6,5,1,16–18]. A peer chooses to remove a partnership if its computed reputation score falls below a locally defined reputation threshold. Eventually, content polluters are identified and isolated from all other peers, and stop receiving the live content.

We also designed a dynamic threshold mechanism that allows previously detected polluters to rehabilitate themselves and rejoin the live transmission. The mechanism works by letting each peer independently change its local reputation threshold depending on the status of the system as perceived by it. If the peer senses the P2P system is currently free of attacks, it lowers its threshold allowing new partnerships with lowly reputed peers, thus giving them a chance to regain their reputations. If, otherwise, the peer detects that the system is currently under attack, it raises its threshold to identify and penalize polluters more quickly.

We evaluated SimplyRep, with both simulation and experiments in a real setup running on PlanetLab [19], comparing it against two previously proposed approaches to deal with malicious peers, namely, a centralized black list and a distributed reputation system called StRepS [2], in various scenarios. Recall that SimplyRep relies only on the individual experiences of the peer to compute reputations. StRepS, in contrast and like various previous reputation mechanisms [9–12], computes reputations by using the individual experiences of the local peer and of its partners (referred to as network testimony), combining them into a single reputation score. This difference makes it a good baseline to evaluate our new method. Our evaluation was performed with varying numbers of polluters as well as with and without collusion of content polluters. Moreover, we also evaluated the impact of key parameters of SimplyRep on its effectiveness, as well as its robustness to dissimulation and whitewashing attacks. In the former, content polluters dynamically change their behavior, by alternating between sending polluted content and forwarding only legitimate content. In the latter, polluters repeatedly leave and rejoin the system with new identities, aiming at loosing their prior reputations.

We highlight three main results in this work. First, to motivate the need of a method to detect content polluters, we show that simply detecting and discarding polluted content before forwarding it is not an effective defense strategy as polluters remain active flooding the system with polluted content. Moreover, peers have to request new copies of the polluted data received, generating significant data retransmission overhead and/or unacceptable delays and data loss in the transmission. Second, both black listing and StRepS do not achieve good results in case of a collusion attack from a reasonably large number of polluters, with overheads due to data retransmission of at least 90% and 30%, respectively. Specifically, we found that, by relying on the network testimony to build reputation scores, StRepS becomes very vulnerable to divergences among the individual experiences. In contrast, the new SimplyRep is able to identify and isolate polluters very fastly, even under collusion and dissimulation attacks, presenting a retransmission overhead that quickly drops to less than 2%. Finally, we show that SimplyRep is reasonably robust to the challenging whitewashing attack, presenting an overhead that, despite somewhat larger than in the other scenarios considered, is still reasonably low (8%) after an initial convergence period.

The rest of this paper is organized as follows. Section 2 presents basic concepts of P2P live streaming systems and discusses the impact of content pollution on these applications. Section 3 introduces our new decentralized reputation mechanism to fight content pollution, and also describes two previous strategies that are here compared against our method. Our evaluation methodology is discussed in Section 4, and our main results are presented in Section 5. Section 6 reviews other related studies. Finally, Section 7 concludes the paper.

2. Content pollution in P2P live streaming systems

In this section, we review the fundamental concepts and mechanisms adopted by currently popular P2P live streaming systems, on which our simulation and PlanetLab experiments are based (Section 2.1). We also report on the results of an experiment to assess the impact of a pollution attack against a currently very popular P2P live application (Section 2.2).

2.1. Live P2P streaming: basic concepts

Various currently popular P2P live streaming systems, such as SopCast, PPLive and GridMedia,[2] use a non-

[2] www.sopcast.com, www.pplive.com and www.gridmedia.com.cn, respectively.