Computer Networks 60 (2014) 115–128

Contents lists available at ScienceDirect

Computer Networks

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

Delay-based congestion control: Flow vs. BitTorrent swarm perspectives

Claudio Testa, Dario Rossi *

Telecom ParisTech, 46 Rue Barrault, 75013 Paris, France

ARTICLE INFO

Article history:
Received 30 April 2013
Received in revised form 18 December 2013
Accepted 19 December 2013
Available online 25 December 2013

Keywords:
BitTorrent
Congestion control
Lower-than best effort congestion control
LEDBAT
TCP-LP
NICE

ABSTRACT

BitTorrent, one of the most widespread file-sharing P2P applications, recently introduced LEDBAT, a novel congestion control protocol aiming at (i) limiting the additional delay due to queuing, to reduce interference with the rest of user traffic (e.g., Web, VoIP and gaming) sharing the same access bottleneck, and (ii) efficiently using the available link capacity, to provide users with good BitTorrent performance at the same time.

In this work, we adopt two complementary perspectives: namely, a flow viewpoint to assess the Quality of Service (QoS) as in classic congestion control studies, and a BitTorrent swarm viewpoint to assess peer-to-peer users Quality of Experience (QoE). We additionally point out that congestion control literature is rich of protocols, such as VEGAS, LP, and NICE sharing similarities with LEDBAT, that is therefore mandatory to consider in the analysis. Hence, adopting the above viewpoints we both (i) contrast LEDBAT to the other protocols and (ii) provide deep understanding of the novel protocol and its implication on QoS and QoE.

Our simulation based investigation yields several insights. At flow-level, we gather LEDBAT to be lowest priority among all protocols, which follows from its design that strives to explicitly bound the queuing delay at the bottleneck link to a maximum *target* value. At the same time, we see that this very same protocol parameter can be exploited by adversaries, that can set a *higher target* to gain an unfair advantage over competitors. Interestingly, swarm-level performance exhibit an opposite trade-off, with *smaller targets* being more advantageous for QoE of BitTorrent users. This can be explained with the fact that larger delay targets slow down BitTorrent signaling task, with possibly negative effect on the swarming protocol efficiency. Additionally, we see that for the above reason, in heterogeneous swarms, any delay-based protocol (i.e., not only LEDBAT but also VEGAS or NICE) can yield a competitive QoE advantage over loss-based TCP.

Overall this tension between swarm and flow-levels suggests that, at least in current ADSL/cable access bottleneck scenarios, a safe LEDBAT operational point may be used in practice. At the same time, our results also point out that benefits similar to LEDBAT can also be gathered with other delay-based protocols such as VEGAS or NICE.

1. Introduction

Pioneered by Jain [22] in late 80 s, delay-based Congestion Control (CC) has been out for a long time, with notable proposals over the years such as VEGAS [8] in late 90s, NICE [42] and LP [25] in early 2000 and more recently LEDBAT [35] in 2010.

The idea of this branch of protocols is to use the variation in the end-to-end delay transmission as *early* congestion signal: in other words, a growing delay beyond a baseline is interpreted as queuing delay building up at

* Corresponding author. Tel.: +33 145817563.
E-mail address: dario.rossi@enst.fr (D. Rossi).

1389-1286/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.bjp.2013.12.018

the bottleneck link, and the amount of data to be sent at every time frame is updated accordingly. This design choice is orthogonal to the one adopted by loss-based protocols, such as in classic TCP NewReno [18] which instead uses packet loss as a *late* congestion signal to tune the data transmission.

Since loss-based protocols forcibly fill the buffer, this can translate into rather large delays, especially at the access link where buffer sizes are relatively large compared to the narrow capacity of ADSL and cable modems. As recent work pointed out, is not uncommon that queuing delays exceed the Earth-to-Moon propagation delay [23,13], for which the "bufferbloat" term was recently coined [12].

Clearly, such huge delays can harm the Quality of Experience (QoE) of interactive communication – including Voice over IP (VoIP), gaming and Web browsing. Additionally, since the bottleneck is placed at the user access link, this means that the user is self-inflicting this QoE degradation, as his own traffic is competing for the bottleneck resources. In other words, QoE degradation results from sustained uploads carried on TCP, whose loss-based Additive Increase Multiplicative Decrease (AIMD) protocol forces the buffer to fill prior to halve the congestion window due to losses.

It follows that bufferbloat can be induced by any application transferring large data volumes over TCP, such as any upload to the Cloud (e.g., Picasa, DropBox, Flickr, etc.), or peer-to-peer file-sharing (e.g., BitTorrent, eDonkey, etc.). To avoid harming contemporary interactive communication of the same user, application developers have thus the choice to exploit alternatives to the standard loss-based TCP behavior. This is precisely the choice of BitTorrent, that recently replaced loss-based TCP with delay-based LEDBAT for data transfer.

This evolution motivates our first viewpoint. As the new protocol is used in BitTorrent swarms, it is important to assess its impact on the quality of BitTorrent users experience – mainly, their completion time [26]. Interestingly though, the protocol has been normalized at the IETF under the Low Extra Delay Background Transport (LEDBAT) in late 2012 [36]. This motivates the second viewpoint: as the protocol is normalized at the IETF, its scope is wider than the BitTorrent ecosystem, and its impact on other applications has to be assessed as well.

In this work, we investigate LEDBAT by means of verb `ns2` simulations, and compare it to other delay-based protocols such as LP, VEGAS and NICE, from both flow vs. swarm perspective. Moreover, in case of LEDBAT we carry out a sensitivity analysis over its main parameter, namely the queuing delay target, to assess the impact of heterogeneous settings. This is an important study, since the parameter can be easily modified by legitimate end-users or legacy implementations (complying to the RFC specification) or by malicious users and developers (violating RFC specification) to possibly gain an unfair advantage. At flow-level, we study the Quality of Service (QoS) of backlogged flows, expressed as the usual network-centric metrics of congestion control studies, such as link efficiency, throughput, and packet loss. At swarm-level, we instead study the Quality of Experience (QoE) of BitTorrent users, expressed as the torrent completion time, that collectively depends on the performance of multiple flows, in a furthermore non-trivial way as we shall see.

Summarizing our most interesting findings, we have that heterogeneous LEDBAT target settings yield to significant unfairness, which is especially true for backlogged connections, where flows with slightly *higher delay target* can starve competing flows. Interestingly though, competitive advantage for selfish users in the swarm case are obtained for *lower delay target* – which suggests that safe LEDBAT operational points may be used in practical cases. At the same time, our results also point out that benefits similar to LEDBAT can also be obtained with other delay-based protocols such as VEGAS or NICE.

The remainder of this work is organized as follows. Related work are discussed in Section 2, while a detailed overview about the congestion control protocols we consider in this study is reported in Section 3. Flow vs. swarm perspectives are then adopted in Section 4 vs. Section 5 respectively. For both perspectives, we investigate the novel LEDBAT protocol (e.g., carrying out a detailed sensitivity analysis of the queuing delay target parameter, and especially of heterogeneous target settings) and contrast performance with that achievable under LP, VEGAS or NICE. Finally, our findings are summarized and discussed in Section 6.

2. Related work

While our most important findings arise in the opposite implications of LEDBAT target settings in the flow vs. swarm perspectives, we point out that, so far, all related effort has focused on either viewpoint in isolation. Hence, we separately treat the above perspectives in this section. In more details, at flow-level, we overview delay-based congestion control protocols and focus on recent work targeting LEDBAT. At swarm-level, we overview studies of BitTorrent performance, and focus on work targeting the impact of packet-level dynamics on content distribution performance.

2.1. Flow viewpoint

Congestion control is a long studied subject: as it would be out-of-scope to provide a full review of the existing literature here, we concentrate on the subset that is most relevant for our work. We present four different categories of

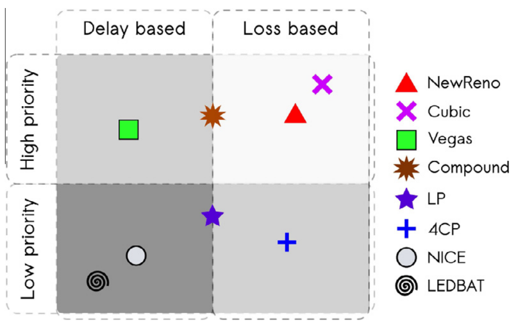

Fig. 1. Congestion control design space: aggressiveness vs. design strategy.