

Evaluating Application-Layer Traffic Optimization Cost Metrics for P2P Multimedia Streaming

Poderys, Justas; Soler, José

Published in: Proceedings of 25th Telecommunications forum

Link to article, DOI: 10.1109/TELFOR.2017.8249299

Publication date: 2017

Document Version Peer reviewed version

Link back to DTU Orbit

Citation (APA): Poderys, J., & Soler, J. (2017). Evaluating Application-Layer Traffic Optimization Cost Metrics for P2P Multimedia Streaming. In *Proceedings of 25th Telecommunications forum* IEEE. https://doi.org/10.1109/TELFOR.2017.8249299

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

• Users may download and print one copy of any publication from the public portal for the purpose of private study or research.

- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Evaluating Application-Layer Traffic Optimization Cost Metrics for P2P Multimedia Streaming

Justas Poderys, Student Member, IEEE, and Jose Soler, Member, IEEE

Abstract — To help users of P2P communication systems perform better-than-random selection of communication peers, Internet Engineering Task Force standardized the Application Layer Traffic Optimization (ALTO) protocol. The ALTO provided data-routing cost metric, can be used to rank peers in P2P communication systems. However, the method to derive the data-routing cost metric is undefined by the standard. This paper proposes and evaluates three methods to derive the data-routing cost metric for use in ALTO servers.

Keywords — Traffic Engineering, Multimedia Streaming, Peer-to-Peer Communication.

I. INTRODUCTION

CCORDING to an industry forecast [1], Internet video will be the second fastest growing category of IP traffic and will make up to 80% of all consumer Internet traffic by 2020. In 2016, 50% of mobile traffic in the Americas, Asia, and Europe was due to video streaming [2]. One of the viable methods to reduce the load on multimedia servers and content delivery networks is to spread the communications load by employing peer-topeer (P2P) assistance schemes. This method is well known and is already widely used in systems such as BitTorrent [3]. While there is already a number of protocol proposals for P2P-assisted multimedia streaming [4], deploying such systems for widespread use is not trivial. To provide users with high Quality-of-Experience (QoE), P2P-assisted multimedia streaming systems require timely data delivery, which in turn requires fast data connections between the peers. Connections between the peers should have enough available bandwidth, low packet-loss rate, latency and jitter [5]. However, in P2P-assisted multimedia streaming systems, there is no way for a peer to know the parameters of the data links to other peers a priori.

To help users of P2P-assisted communication systems with the selection of communication peers, the Internet Engineering Task Force (IETF) standardized the Application-layer Traffic Optimization (ALTO) protocol [6]. By using the ALTO protocol together with an ALTO server, peers in a P2P-assisted system can discover the

network topology and use this information to rank potential communication peers. Unlike some other previously proposed protocols providing similar service [7]–[9], ALTO is intended to be deployed by the ISPs. Such deployment model allows the ALTO server to have better information about the underlying network and its parameters (for example topology, routing data, links load). While the ALTO protocol specification [6] defines the ALTO information model and the format of the data request and response messages, it leaves the method to derive the datarouting cost-metric undefined.

This paper proposes three data-routing cost-metrics for use in ALTO servers. To test how these metrics impact the QoE of P2P-assisted multimedia streaming, a reference network model was created and emulated using the CORE computer network emulator [10]. The performed tests compare how the startup time and playback continuity are affected by the number of P2P users in the network, background data-traffic levels and method of deriving the ALTO cost-metric. In addition to testing the QoE parameters, the tests also show how using ALTO can localize the data exchanges between the users and reduce the load on the multimedia streaming servers.

II. RELATED WORK

A substantial part of the global Internet traffic is generated by P2P-based data exchange [1]. A big part of previous work in the area [11]–[14] derives the ALTO cost metric from the global Internet topology. This is achieved by counting the number of Autonomous Systems (AS) that are traversed (AS-path length) when data travels between the source and destination addresses.

While the above mentioned papers indicate that such approach reduces data download times and localizes data traffic to the originating-AS, it has several shortcomings. By using the AS-path length as a cost metric, ALTO server hides information about the internals of the network. By doing so, highly congested networks are treated the same way as congestion free networks. Furthermore, using the AS-path length as a peer selection metric does not require an ALTO server. All nodes connected to the Internet can perform a path trace to the remote peer, and calculate the number of AS that the path crosses.

III. ALTO PROTOCOL COST METRICS

The ALTO Problem Statement [15] describes several

Justas Poderys is with the Department of Photonics Engineering, Technical University of Denmark, Oersted pl. 343, 2800 Kgs. Lyngby, Denmark (phone +45 45 25 36 13; e-mail: juspo@fotonik.dtu.dk).

Jose Soler is with the Department of Photonics Engineering, Technical University of Denmark, Oersted pl. 343, 2800 Kgs. Lyngby, Denmark (phone +45 45 25 38 08; e-mail: joss@fotonik.dtu.dk).

use cases for the ALTO protocol. In all use cases, ALTO services are used to provide applications with information about the underlying network's topology, in order to perform better-than-random selection of communication peers. This selection is made possible by assigning a cost value to send data from the source peer to the destination peer. If several network peers can serve the same data, then the ALTO cost metric for each source-destination addresses pair (where destination is the node making an ALTO query) can be used to select the best source peer.

As described in the related work section, several previous works in the field used the AS-path length as the ALTO cost metric. To overcome some of the limitation of using AS-path length as a cost metric, this paper proposes three methods to derive the ALTO data routing cost (referred as cost metrics). The following section describes each cost metric and compare them according to the following 4 criteria:

Relative accuracy. The cost metric should reflect the network's topology and load conditions. The metric should also change, when the network state changes.

Calculation complexity. The cost metric should be fast to calculate to allow short ALTO server response times.

Cacheability. The cost metric, once calculated, should be stored in server's memory to prevent recalculation.

Interoperability. The cost metric should be calculated in a way, that allows a composite metric (metric for data path crossing several networks) to be calculated.

This work proposes these three cost-metrics: Hops-Routing-Cost (HRC); Ospf-Routing-Cost (ORC) and Path-Residual-Bandwidth (PRB).

The first analyzed metric derives the data-routing cost value from the number of routers along the data-path. This approach follows the routing metric used in the RIP routing protocol [16], albeit without the 15 hops limit. The routing cost metric is increased by 1 for each router that the data passes between the source and destination addresses. This method is referred as Hops-Routing-Cost (HRC). The second metric derives the cost value from the Interior Gateway Protocol (IGP) used in the ISP's network. Experiments in this work use Open Shortest-Path First (OSPF) as an IGP protocol, and the cost value is calculated by following the OSPF protocol specification [17]. This metric is referred as Ospf-Routing-Cost (ORC). The third metric derives its cost value from the end-to-end available bandwidth along the data-path [18]. The value of the cost metric is equal to the smallest available single link's bandwidth in the links that data traverses along the datapath and is referred as Path-Residual-Bandwidth (PRB). For the first two metrics, a path with a lower metric value is preferred to a path with a higher value. For the PRB metric, the relation is the opposite: a path having a higher value is preferred to a path having a lower value.

Among the three metrics, PRB most accurately reflects the current network conditions, as it is derived directly from the network link load values. The ORC metric is less accurate, as it takes into account only the provisioned link capacities and not the actual load. The HRC is the least accurate, because it indicates only the number of crossed routers.

When an ALTO server is deployed in networks with simple topologies, the ORC metric is the easiest to calculate. It can be done by looking up the value in the first router having a full routing table (routes to all network prefixes) along the data path. In complex networks (networks with multiple OSPF areas), the ALTO server will have to trace the complete path between the source and destination addresses. This is due to the fact that routers have full view only of the OSPF area they are operating in. Deriving the HRC value is more computationally demanding, as it requires the ALTO server to always trace the complete data path from the source address to the destination address. When the ALTO server has routing data from all routers in the network, the HRC value can be calculated in linear time, because the shortest-path is already computed by the routers. Calculating the PRB requires most resources: in addition to tracing the data path between the source and destination addresses, the ALTO server has to lookup the provisioned link capacities and calculate the average links' load.

Caching of the calculated data-routing cost values is an important consideration in ALTO servers handling high number of user queries. An operator of an ALTO server might want to re-use the calculated cost value, when the data path is between the PIDs, for which cost was calculated recently. Once calculated, the ORC and HRC values can be cached for as long as there are no network topology changes and loadbased routing is not used. However, the PRB cost value is valid only for a short time. The value changes when the next observation of network links utilization is completed.

The current version of the ALTO protocol does not define the interface between ALTO servers running in different ISP networks [6]. At the same time, it is worth to consider, how each of the different cost metrics would interoperate, once the interface is standardized. The HRC metric is the easiest to interoperate with, since it is derived strictly from the network's topology. When a data-path crosses several ISP networks, the composite metric is the sum of individual metrics from each ISP's network. In a similar manner, the PRB composite metric is equal to the minimal value of all metrics in each ISP network. Of all the three metrics, the ORC is the least interoperable due to two reasons. First, not all networks use OSPF as their IGP, and other widespread IGPs (e.g., IS-IS, EIGRP) have incompatible cost metrics. Second, the cost value of each OSPF link is derived as a ratio with the reference bandwidth [17].

IV. EXPERIMENTAL SETUP

Tests to evaluate the different cost metrics were carried out using the CORE network emulator [10]. The test network topology was adapted from an industry's whitepaper [19]. The test network contains 96 user-nodes divided into 16 groups (PIDs) of 6 nodes. In the tests, usernodes are connected to the network using Digital Subscriber Line (DSL) technology. This is done by connecting each group of users to a Digital Subscriber Line Access Multiplexer (DSLAM), which in turn is connected to a Broadband Network Gateway (BNG). Every BNG is connected to two DSLAMS, and all BNGs are connected in a ring topology. Links between the users and DSLAMs were provisioned with 10 Mbps capacity, and links between the DSLAMs and BNGs were provisioned with 30 Mbps capacity, giving the DSLAMs an oversubscription ratio of 1:2.

The impact of ALTO on users' QoE is evaluated by observing two objective multimedia streaming quality parameters [20]: start-up delay and playback continuity index. In addition to the parameters impacting user's experience, this work also considers data traffic localization and the multimedia server's load reduction. In previous works using the ASpath length as a cost metric [11]–[14], one of the goals of using ALTO was to localize the data traffic localization to the BNG connecting a user to the rest of the network.

V. EXPERIMENTAL RESULTS

A. Start-up Time

The comparison is started by measuring the average startup time of P2P clients when using different cost metrics. The observed average start-up time based on the number of users and background data-traffic levels is show in Fig. 1. In both use cases, lines indicate the average startup time using each cost-metrics at the indicated traffic level.



Fig. 1. Average start-up time in Live and VoD use-cases with different background data traffic levels.

Fig. 1 shows that the start-up time increases together with the increasing number of users and background traffic levels. However, start-up time is not affected by the costmetric, as lines representing each individual metric are all clustered together. The reason for this is two-fold. First, peers in P2P systems make requests to download data based on the data availability. Only if several peers have the same data, can then they be ranked with the help of ALTO. In this case, as soon as the multimedia server indicates that it has new data available, all users request this data, regardless of the cost metric assigned to the source node. The second cause of the indifference in the cost metric is the design of the experiments. In the experiments, all nodes running P2P software started at the same time (flash-crowd arrival). The reasons for choosing this method of nodes arrival are two. First, it is the hardest use-case for P2P streaming systems and hence the results show worst-case scenarios. Second, such nodes arrival pattern is experienced by content delivery networks during high demand periods, such as beginning of sport events or release of popular TV-series.

B. Playback Continuity

Next, this work considers how using different cost metrics impact the Playback Continuity Index (PCI) of multimedia in P2P clients. The reference "No ALTO" use-case is shown in Fig. 2a. The remaining three sub-figures each depict PCI for different cost metrics.



Fig. 2. Playback Continuity Index (PCI) for each costmetric based on number of users, usage scenario and background traffic levels.

The comparison is started with the VoD use-case. In the reference "No ALTO" case, the playback index gradually decreases, as the number of users in the network increases. The HRC and ORC metrics perform very similarly, with the playback continuity index above 0.65 for all numbers of users. The PRB performs worse, but still better than the reference "No ALTO" case. It is important to note, that the HRC and ORC metrics perform equally well in all background traffic levels, while the performance of PRB method decreases, as the background data-traffic levels increase. Continuing with the Live use case, it can be seen from Fig. 2, that the different cost metrics have little impact on the playback continuity index.

C. Data Traffic Localization

Next, this paper considers how ALTO impacts data transmissions localization. Here localization means influencing users to exchange P2P data with other users connected to the same BNG. Localization of data transmissions is important to ISPs, because it reduces the amount of data that traverses the operator network. Fig. 3 shows the average number of data messages received by each P2P user based on the origin of data pieces, and usage scenario for each cost metric. Here, "Local" refers to the data pieces received from the users connected to the

same BNG, and "Remote" - to data pieces received from users connected to other BNGs and not the streaming server.

Fig. 3 shows, that in both Live and VoD scenarios, using the ORC and HRC metrics increase the number of data pieces received from the users connected to the same BNG. However, the PRB metric is not as effective at localizing data traffic. Comparing the VoD (Fig. 3a) and Live (Fig. 3b) use cases, it is important to note the different vertical scales. Taking the HRC metric as an example: each user received on average 373 data pieces from other users connected to the same BNG in Live use case, compared to 2314 data pieces in VoD use case. This further reinforces our argument, that in Live usecase, peers receive most of the data from the streaming server, reducing the utility of ALTO.



Fig. 3. Average number of data messages received by users based on the communications peer location, use-case and cost metric.

VI. CONCLUSIONS AND FUTURE WORK

This work proposes three methods (HRC, ORC and PRB) to derive the ALTO data-routing cost metric. The proposed methods were implemented in an ALTO server and tested in an emulated network with up to 50 users performing P2P-assisted multimedia streaming. The experiments show that in the VoD use-case, HRC and ORC methods performed best. Using these methods, users in the test network experienced playback continuity index above 0.65 for all network sizes.

However, among the two methods, HRC is preferred, due to being independent of the IGP used in the ISP network and allowing easier interoperability. While the experiments presented here indicate, that the HRC is the preferred method for deriving the data-routing cost metric, several questions remain open. In this work, the experiments were based on a single network topology, with all P2P users arriving at the start of each experiment: In a future work, we plan to test the impact of ALTO in a wider range of network topologies and with different user arrival and departure patterns. The next open question is which methods can be used to increase the QoE for users streaming Live multimedia.

REFERENCES

- "Cisco visual networking index: Forecast and methodology 2015-2020," Tech. Rep., 2016.
- [2] "Ericsson mobility report western europe," Tech. Rep., 2016.
- [3] B. Cohen, "The bittorrent protocol specification," Jan. 2008. [Online]. Available: http://www.bittorrent.org/beps/bep0003.html

- [4] X. Zhang and H. Hassanein, "A survey of peer-to-peer live video streaming schemes - An algorithmic perspective," Computer Networks, vol. 56, no. 15, pp. 3548–3579, 2012.
- [5] F. Tommasi, V. De Luca, and C. Melle, "Packet losses and objective video quality metrics in H.264 video streaming," Journal of Visual Communication and Image Representation, vol. 27, pp. 7–27, 2015.
- [6] R. Alimi, R. Penno, Y. Yang, S. Kiesel, S. Previdi, W. Roome, S. Shalunov, and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol," RFC 7285, Sep. 2014.
- [7] T. Ng and H. Z. H. Zhang, "Predicting Internet network distance with coordinates-based approaches," Proceedings. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies, vol. 1, no. c, pp. 170–179, 2002.
- [8] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi : A Decentralized Network Coordinate System," Communication, vol. 34, pp. 15–26, 2004.
- [9] B. Wong, A. Slivkins, and E. G. E. E. G. Sirer, "Meridian: A lightweight network location service without virtual coordinates," ACM SIGCOMM Computer Communication Review, vol. 35, no. 4, pp. 85–96, 2005.
- [10] J. Ahrenholz, "Comparison of CORE network emulation platforms," Proceedings - IEEE Military Communications Conference MILCOM, pp. 166–171, 2010.
- [11] J. Seedorf, S. Niccolini, M. Stiemerling, E. Ferranti, and R. Winter, "Quantifying operational cost-savings through ALTO-guidance for P2P live streaming," Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 6236 LNCS, pp. 14–26, 2010.
- [12] L. Guanxiu, Y. Suqi, and H. Xinli, "A Novel ALTO Scheme for BitTorrent-Like P2P File Sharing Systems," 2013 Third International Conference on Intelligent System Design and Engineering Applications, pp. 135–139, 2013.
- [13] D. Alex, L. Perez, S. Henrique, and B. Brito, "Delivering Application-Layer Traffic Optimization Services based on Public Routing Data at Internet eXchange Points."
- [14] N. Pinthong and W. Lilakiatsakun, "Performance of BitTorrent-Like P2P File Sharing Systems inspired by ALTO," IEEE Region 10 Annual International Conference, Proceedings/TENCON, 2013.
- [15] J. Seedorf and E. Burger, "Application-Layer Traffic Optimization (ALTO) Problem Statement," RFC 5693, Oct. 2009.
- [16] G. Malkin, "RIP Version 2," RFC 2453, Nov. 1998.
- [17] J. Moy, "OSPF Version 2," RFC 2328, Apr. 1998.
- [18] M. Jain and C. Dovrolis, "End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput," IEEE/ACM Transactions on Networking, vol. 11, no. 4, pp. 537–549, 2003.
- [19] Network Configuration Example Configuring MX Series Universal Edge Routers for Service Convergence, 2016.
- [20] J. Kuo, C. Shih, C. Ho, and Y. CHen, "Advanced bootstrap and adjusted bandwidth for content distribution in peer-to-peer live streaming," Peerto-Peer Networking and Applications, vol. 8, no. 3, pp. 414–431, 2015.