Poster Abstract: Fair and Efficient Dynamic Bandwidth Allocation with OpenFlow^{*}

Maryam Elahi melahi@mtroyal.ca Mount Royal University, Calgary, AB, Canada Joel van Egmond, Mea Wang, Carey Williamson joel.vanegmond,meawang,cwill@ucalgary.ca University of Calgary, AB, Canada Jean-François Amiot jfamiot@cybera.ca Cybera Inc. Calgary, AB, Canada

ABSTRACT

Large-scale not-for-profit Internet Service Providers (ISPs), such as National Research and Education Networks (NRENs) often have significant amounts of underutilized bandwidth because they provision their network capacity for the rare event that all clients utilize their purchased bandwidth. However, traffic policers are still applied to enforce committed purchase rates and avoid congestion. We present the design and initial evaluation of an SDN/OpenFlow solution that maximizes the network link utilization by user-defined fair allocation of spare bandwidth, while guaranteeing minimum bandwidth for each client.

CCS CONCEPTS

• Networks \rightarrow Network resources allocation; Network performance evaluation.

KEYWORDS

Dynamic Bandwidth Allocation, OpenFlow, SDN, Utilization, Fairness, TCP/IP Performance Evaluation

ACM Reference Format:

Maryam Elahi, Joel van Egmond, Mea Wang, Carey Williamson, and Jean-François Amiot. 2020. Poster Abstract: Fair and Efficient Dynamic Bandwidth Allocation with OpenFlow. In ACM/SPEC International Conference on Performance Engineering Companion (ICPE '20 Companion), April 20–24, 2020, Edmonton, AB, Canada. ACM, New York, NY, USA, 2 pages. https: //doi.org/10.1145/3375555.3383587

1 INTRODUCTION

Internet Service Providers (ISPs) commonly apply Quality of Service (QoS) technology to police or shape client traffic [4]. Commercial ISPs focus on maximizing profits by overselling their channel, with the assumption that not all clients will be using their maximum purchased rate simultaneously. Consequently, the industry's status quo is to sell an "up to" bandwidth rate. In contrast, large-scale not-for-profit ISPs, such as NRENs opt to use QoS controls to provide minimum rate guarantees to their clients, and therefore have significant amounts of underutilized bandwidth [7].

Although there are many mechanisms currently available for capping or limiting bandwidth at a certain level [4], we know of

ICPE '20 Companion, April 20-24, 2020, Edmonton, AB, Canada

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7109-4/20/04.

https://doi.org/10.1145/3375555.3383587

no dynamic and scalable solution that provides minimum bandwidth guarantees to clients, while maximizing utilization of the available bandwidth with fair allocation of the spare bandwidth, where the fair allocation is defined by the ISP administration. The QoS architectures such as IntServ [3] and DiffServ [2] are either too complex and not scalable, or do not provide strong fairness guarantees. For example, the two-rate three-color marker policers guarantee minimum bandwidth allocation, while allowing competing flows to utilize the spare bandwidth. However, they do not support user-specified fair allocation of the spare bandwidth [2].

Software-Defined Networking (SDN) is an emerging paradigm that decouples network control from forwarding functions, enabling network control to be directly programmable. SDN introduces new possibilities for dynamic bandwidth allocation (DBA). OpenFlow, which is the prevalent open standard protocol for SDN, facilitates DBA at line rate with features like per-flow metering and flow pipeline [6]. Previous studies have leveraged SDN capabilities for optimal bandwidth allocation in applications such as stream analytics [1], and end-to-end QoS guarantees [5]. However, fair and efficient DBA with minimum guarantees has not been considered.

We present a DBA system using OpenFlow, which maximizes the network link utilization by fair allocation of spare bandwidth, while ensuring minimum bandwidth guarantees for each client.

2 FLOW AGGREGATION AND METERING

Our DBA solution is implemented using an OpenFlow enabled switch (OpenFlow 1.3 or higher) with flow pipeline, ingress meters, and egress QoS queues. Figure 1 shows the three-tier flow aggregation and metering for our fair and efficient DBA solution. The egress QoS queues follow a strict priority system, in which the higher priority queue always transmit first. The highest priority queue is used for guaranteed traffic, and the lower priority queues are used for traffic over the committed minimum rate based on the dynamic allocation algorithm.

The first tier (flow table 0) aggregates the traffic from each client, and mark the portion within the minimum guaranteed rate with high priority. The high priority traffic is not metered at the second tier. At the third tier (flow table 2), the high priority traffic from all clients are aggregated and directed to the high priority egress QoS queue. This redirection guarantees the minimum rate for each client, as long as the channel is not overbooked.

Traffic exceeding the guaranteed rate is metered at the second tier (flow table 1). The DBA algorithm sets the metering rate at the second tier for each client. Meters at this tier mark traffic within the dynamically allocated maximum rate with medium priority, and traffic beyond this rate with low priority. The marked traffic from all clients is then aggregated at the third tier according to the priority level, and directed to the respective egress QoS queue.

^{*}This work was supported by Mitacs and Cybera through the Mitacs Elevate program.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).



Figure 1: Three-tier flow aggregation and metering for fair and efficient dynamic bandwidth allocation

 Table 1: Results for bandwidth utilization of competing TCP traffic from two clients. (Link capacity 600, rates in Mbps)

Algorithm	Client	Min-Rate	Demand	Fair Share	Delivered
Max-Min Fair	H1	100	400	250	85
Exact-Allocation	H2	200	600	350	519
Max-Min Fair	H1	100	400	250	234
Over-Allocation	H2	200	600	350	324

3 DBA FAIR UTILIZATION

The DBA algorithm is implemented within the SDN controller. The controller monitors the real-time bandwidth usage for each client through the flow statistics at each tier, and dynamically adjusts the metering rate at the second tier in order to maximize utilization. Different allocation algorithms are possible through this design. We have implemented a max-min fair allocation algorithm, which gives an equal share of the spare bandwidth to all clients upon demand, regardless of their committed purchased rate.

There are subtleties in how the spare bandwidth is allocated. With exact-allocation, where the sum of all allocated bandwidth is no more than the link capacity, either the link remains underutilized (when allocation is based on equal share of the spare bandwidth), or the TCP congestion control for some clients adapts to a lower rate than their fair share (when allocation is based on the perceived demand by the controller). To achieve maximum utilization and maintain max-min fairness, we use an over-allocation scheme, which allocates the equal share to clients with low perceived demand (even though they may not utilize it), and allocates the fair share of the unutilized spare bandwidth to clients with high demand. The over-allocation stops when traffic reaches the fair allocation rate for all clients.

4 IMPLEMENTATION AND RESULTS

We have implemented a proof-of-concept testbed using SDN enabled Pica8 (P-3290) switches and Ryu SDN controller with REST API (source code available on GitHub [8]). These switches only implement one stage of the OpenFlow pipeline, and therefore do not support pipeline processing instructions such as goto_table. Therefore, we used three switches to implement the three-tier flow aggregation and metering.

Table 1 shows utilization results for competing TCP traffic. With exact-allocation based on perceived demand, TCP congestion control for the client with the lower min guaranteed rate adapts to a lower rate than its fair share, and the client with higher min rate gets more than its fair share. With over-allocation, both clients get close to their fair share. Figure 2 illustrates the adaptivity of our



Figure 2: Adaptivity of the max-min fair DBA algorithm over time. Client H1 with guaranteed rate 100 Mbps tries to transmit at rate 400 Mbps for 30s. H1 initially utilizes close to 400 Mbps through the link without competition. Client H2 with guaranteed rate 200 Mbps starts competing for bandwidth with 10s delay, and tries to transmit at rate 600 Mbps for 30s. The algorithm quickly adjusts the rate for H1 and H2. When H1 stops transmission, H2 is able to utilize the available spare bandwidth and transmit at near 600 Mbps.

allocation scheme over time. The TCP congestion control quickly adapts to the max-min fair rate when over-allocation is applied.

5 CONCLUSIONS

We present a three-tier flow-aggregation and metering design using OpenFlow that provides fair and efficient DBA for networks with minimum bandwidth guarantees. Initial evaluation on an SDN testbed shows that to achieve max-min fair utilization of the spare bandwidth by TCP flows, an over-allocation scheme is required. Our ongoing work include investigating the interplay between the TCP congestion control and other fair allocation algorithms (e.g. proportional fair), and evaluating the scalability of our solution.

REFERENCES

- W. Aljoby, X. Wang, T. Fu, and R. Ma. 2018. On SDN-enabled online and dynamic bandwidth allocation for stream analytics. In Proc. of IEEE ICNP. 209–219.
- [2] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. 1998. RFC 2475: An architecture for differentiated services.
- [3] D. Clark, R Braden, and S. Shenker. 1994. RFC 1633: Integrated services in the internet architecture: an overview.
- [4] T. Flach, P. Papageorge, A. Terzis, L. Pedrosa, Y. Cheng, T. Karim, E. Katz-Bassett, and R. Govindan. 2016. An Internet-Wide Analysis of Traffic Policing. In Proc. of ACM SIGCOMM Conference (Florianopolis, Brazil). 468–482.
- [5] H. Krishna, N. van Adrichem, and F. Kuipers. 2016. Providing bandwidth guarantees with OpenFlow. In Proc of SCVT. 1–6.
- [6] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. 2008. OpenFlow: enabling innovation in campus networks. ACM SIGCOMM Computer Communication Review 38, 2 (2008), 69–74.
- [7] S. Nickolay, E. Jung, R. Kettimuthu, and I. Foster. 2018. Bridging the gap between peak and average loads on science networks. *Future Generation Computer Systems* 79 (feb 2018), 169–179.
- [8] J. van Egmond and M. Elahi. 2019. Ryu REST Controller for Fair Dynamic Bandwidth Allocation. https://github.com/bmelahi/ryuRestDBA.git.