Intent-Based Policy Optimization in SD-WAN

Pham Tran Anh Quang, Sebastien Martin, Jeremie Leguay
Huawei Technologies Ltd.
Paris Research Center, France.

Xu Gong, Feng Zeng
Huawei Technologies Ltd.
Dongguan Research Center, China.

ABSTRACT

To optimize bandwidth utilization in wide area networks, a controller typically maintains policies at edge routers. In this context, our demonstration presents a versatile policy optimization model that carefully selects the set of overlay links for each application based on its requirements and the overall intent of the operator. The optimization of policies is realized using an SLA prediction model for several intents. We demonstrate, for instance, that latency is improved by 40% when the high-quality intent is selected.

CCS CONCEPTS

• **Networks** → **Traffic engineering algorithms**.

ACM Reference Format:
Pham Tran Anh Quang, Sebastien Martin, Jeremie Leguay and Xu Gong, Feng Zeng. 2021. Intent-Based Policy Optimization in SD-WAN. In *ACM SIGCOMM 2021 Conference (SIGCOMM '21 Demos and Posters), August 23–27, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3472716.3472858

1 INTRODUCTION

Quality of Service (QoS) and routing mechanisms are key to control how bandwidth is shared among the different applications in Software-Defined Wide Area Networks (SD-WAN) [9]. They can be used to serve multiple *intents* to optimize the utilization of the network, e.g., minimization of financial expenses or maximization of the experienced quality. In typical SD-WAN architectures, a centralized controller maintains a set of policies deployed at edge routers that interconnect multiple sites (e.g., enterprise sites, data centers). Each edge router is configured to send traffic to the others over several access transport networks (e.g., private lines based on MPLS, cheaper broadband Internet). Ingress routers are responsible for the load balancing of flows across outgoing networks so as to satisfy Service Level Agreements (SLA) in terms of QoS, security, etc.

Policy optimization algorithms have been proposed for several purposes. In [3], the authors minimize costs for total volume and 95^{th} percentile charging rules. Google's B4 [5] optimizes the fair sharing of the available bandwidth for elastic applications. To optimize latency and other QoS parameters, a number of works have applied Deep Reinforcement Learning (DRL) [8], under the umbrella of *experience-driven networking* [8]. Closed-form performance models have also been embedded into routing optimization algorithms to minimize latency. For instance, [1] considered the Kleinrock function [6].

Our demonstration showcases a policy optimization model for Smart Policy Routing (SPR) [4] to either minimize cost or latency, while meeting QoS requirements of applications. It decides the set of transport networks, i.e. overlay links, that each application is allowed to use. Our solution goes beyond state of art as it can address multiple intents from the network operator while satisfying individual requirements from applications. Furthermore, to take strategic decision on top of the tactical decisions taken inside routers, our model integrates accurate SLA predictions for latency.

2 INTENT-BASED POLICY OPTIMIZATION

Our architecture is based on a centralized controller that can modify periodically policies at ingress routers.

Smart Policy Routing. Ingress routers are configured with an SPR policy for each application, also called flow group. The policy contains the set of candidate overlay links that a flow group is allowed to use. All the policies are configured, in a slow control loop, by the network controller and executed inside devices, in a fast control loop. At a high frequency, each device determines for each flow group the set of active links from the pool of candidate links authorized by the controller based on link-level measurements (i.e., latency, jitter, loss). The set of active links is used to load balance traffic for each flow group. In order to avoid congestion and satisfy SLA requirements of flow groups, policies have to be carefully decided by the controller to mitigate interferences between flow groups. The main challenge addressed by our demonstration is to automatically configure these policies for several intents.

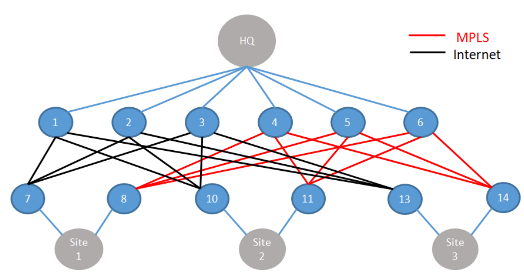

Figure 1: SD-WAN network scenario.

Problem Formulation. We consider an SD-WAN network as shown in Fig. 1 with a set of overlay links E and a set of flow groups K. For each flow group $k \in K$, the measured traffic demand is denoted b^k and the set of all possible overlay links is given by $E_k \subseteq E$ (e.g., outgoing links at ingress routers). Let $x \in [0,1]^{|K| \times |E_k|}$ be the split ratio of traffic on link $e \in E_k$ for flow group $k \in K$ (x_e^k) and $D \in \mathbb{R}_+^{|K|}$ be the maximum delay required by the flow group k (D_k). For a piece-wise modeling of the cost function, let $LU \in [0, C_e^i]^{|E|}$

SIGCOMM '21 Demos and Posters, August 23–27, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-8629-6/21/08.
https://doi.org/10.1145/3472716.3472858

be the link utilization of the i-th piece and C_e^i its capacity. For each link e and each piece i, $w_e^i \in \mathbb{R}_+$ represents its cost per unit of traffic (remark that $w_e^i \leq w_e^{i+1}$). The load balancing policy can be found solving the following optimization model for two intents:

$$\min \quad \sum_{k\in K}\sum_{e\in E_k}\Big[\sum_{i\in I_e} w_e^i LU_e^i \text{ (min cost) or } f_e^k(x) \text{ (high quality)}\Big]$$

$$\text{s.t.} \quad \sum_{i\in I_e} LU_e^i = \sum_{k\in K_e} b_k x_e^k \leq C_e, \quad \forall e \in E, \tag{1}$$

$$f_e^k(x) \leq D_k \quad \forall k \in K, \forall e \in E_k, \tag{2}$$

$$\sum_{e\in E_k} x_e^k = 1, \quad \forall k \in K \tag{3}$$

where $f_e^k(x)$ is the delay function that provides the delay of flow group k on link e by considering the assignment given by x. Constraints (1) ensure the capacity of each link is satisfied. Inequalities (2) compute the delay of each flow group. Constraints (3) ensure that all the traffic demand is routed.

Full model. In our implementation, the SLA prediction f relies on a queuing model for a non-preemptive priority scheduler with 3 classes (real-time, business and bulk). The full optimization model also supports other intents (e.g., MLU minimization), stickiness constraints to limit modifications of policies and a proper estimation of UCMP weights proportionally to link capacities.

3 DEMONSTRATION

We use NS3 simulator [7] with Open Flow 1.3 [2]. Applications are generating traffic following typical patterns for 3 flow groups (Real-time, Business, Bulk) with end-to-end delay requirements of respectively 40ms, 60ms, and 1s. The transport layer is TCP. The microflow inter-arrival time varies to generate diurnal traffic patterns. As Fig. 1 shows, the topology is an SD-WAN network where 3 branches are connected to a headquarter site using 3 MPLS and 3 broadband Internet access networks. The capacity of Internet links is 1.5 times that of MPLS links. The propagation delay of Internet and MPLS links range from 35 ms to 50 ms and 10 ms to 20 ms, respectively. Traffic is prioritized using a non-preemptive priority scheduler. The priority of packets is marked by Differentiated Services Code Points (DSCP). Link-level measurements are collected every 1s and policies are updated by the controller every 5s. In real networks, these periods are expected to be much larger because traffic varies at a slower pace than in this demonstration. Fig. 2 shows the diurnal evolution over time of the total traffic (throughput of all flow groups) and the corresponding end-to-end delays when the high-quality intent is used. Business traffic is 2 times more than real-time and bulk traffic. End-to-end delay of Real-time is much smaller than for Business and Bulk. Business also has a stable and low delay most of the time. There is only one delay spike around 320s due a rapid increase in traffic (faster than policy update). Bulk plots more delay spikes than others especially during high traffic periods due to its low priority. Table 1 provides the end-to-end delay, SLA violations (% of epochs the target is not met), and the total cost for different intents `OPT-HQ` (opt. model with high-quality intent), `OPT-COST` (opt. model with minimum cost intent). We also compare with `All-TNs` where all links can be used and the selection of links is handled by SPR [4]. For the charging model, if traffic is less than a threshold (e.g. 60% of link capacity) the cost is 150$ per traffic unit. Otherwise, it is 10 times more expensive. `OPT-HQ` obtains the lowest delay, up to 40% lower than `All-TNs` for Real-time and Business. Thanks to the estimation of delay in the optimization model, there is reduction in SLA violations, up to 3 times compared to `All-TNs`. With the minimum cost intent, there is no remarkable difference with `All-TNs` in delay, but still a significant gap for SLA violations. In addition, the minimum cost intent yields a cost 12% cheaper. The video is available here: https://tinyurl.com/wwkn237t

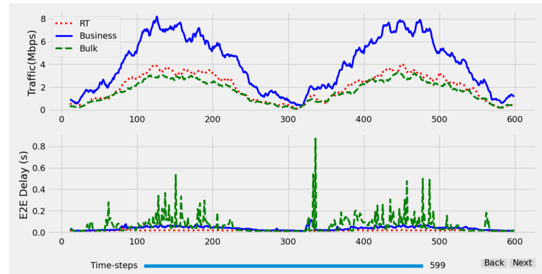

Figure 2: Traffic and end-to-end delay (high-quality intent).

Policy	`OPT-HQ`	`OPT-COST`	`All-TNs`
		Real-time	
Average delay (ms)	16.1	22.2	25.2
SLA violations (% time)	0.0	0.0	3.84
		Business	
Average delay (ms)	36.5	42	43.9
SLA violations (% time)	7.23	11.8	14.4
		Bulk	
Average delay (ms)	61.7	65	63.7
SLA violations (% time)	0.0	0.0	0.7
Total cost	8033	7077	8016

Table 1: End-to-end delay, % SLA violations and total cost

REFERENCES

[1] Walid Ben-Ameur and Adam Ouorou. 2006. Mathematical models of the delay constrained routing problem. *Algorithmic OR* 1, 2 (2006).
[2] Luciano Jerez Chaves, Islene Calciolari Garcia, and Edmundo Roberto Mauro Madeira. 2016. OFSwitch13: Enhancing Ns-3 with OpenFlow 1.3 Support. In *Proc. ACM Workshop on Ns-3*. New York, NY, USA.
[3] Zbigniew Duliński, Rafał Stankiewicz, Grzegorz Rzym, and Piotr Wydrych. 2020. Dynamic traffic management for sd-wan inter-cloud communication. *IEEE Journal on Selected Areas in Communications* 38, 7 (2020), 1335–1351.
[4] Huawei Technologies. 2021. *Smart Policy Routing*. Technical Report. https://support.huawei.com/enterprise/en/doc/EDOC1000174111/2ed71c78/smart-policy-routing
[5] Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, et al. 2013. B4: Experience with a globally-deployed software defined WAN. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 3–14.
[6] Leonard Kleinrock. 2007. *Communication nets: Stochastic message flow and delay.* Courier Corporation.
[7] George F. Riley and Thomas R. Henderson. 2010. *The ns-3 Network Simulator.*
[8] Zhiyuan Xu, Jian Tang, Jingsong Meng, Weiyi Zhang, Yanzhi Wang, Chi Harold Liu, and Dejun Yang. 2018. Experience-driven networking: A deep reinforcement learning based approach. In *IEEE INFOCOM.*
[9] Zhenjie Yang, Yong Cui, Baochun Li, Yadong Liu, and Yi Xu. 2019. Software-defined wide area network (SD-WAN): Architecture, advances and opportunities. In *Proc. of IEEE ICCCN.*