Admission Control with Machine Learning in Software Defined Networks

Jérémie Leguay, Lorenzo Maggi, Moez Draief, Stefano Paris, Symeon Chouvardas,
Huawei Technologies, FRC Research Lab
(Boulogne-Billancourt, France)
Email: firstname.lastname@huawei.com

1. Admission control in SDNs

Software-Defined Networking (SDN) technologies have radically transformed network architectures. They provide programmable data planes that can be configured from a remote controller platform. This creates an opportunity to implement routing processes that are more efficient than classic ones: in fact, the controller can take real-time decisions at a (logically) centralized location using an accurate and global view of the network. Moreover, SDN controllers have a tremendous computational power compared to legacy embedded devices. This encourages the development of smarter network control planes using cutting-edge optimization and machine learning techniques. A key task of the SDN controller is the Admission Control on incoming connection requests, in an online manner. Its goal is to gracefully manage service requests when the network becomes highly utilized, as new incoming requests arrive. It accepts or drops new requests depending on the resource availability. Non-myopic decisions have to be made to maximize a given profit, such as the total accepted throughput, the financial revenue generated or the quality of service experienced by users. Admission Control can be formulated as an online packing problem, as the goal is to maximize the number of accepted requests (subject to feasibility constraints). The challenge in this context comes from the online nature of the optimization problem. New variables and additional constraints are revealed sequentially, as soon as an arrival or departure of a flow occurs in the system. The theory on online algorithms has evolved significantly in the last decade for this type of problems. Algorithms with guaranteed over the offline optimal, which has full information on the future state of the system, have been proposed for online packing problems. For these guarantees to hold, admission decisions need to be consistently taken. The centralized nature of SDN lies the ground to apply online algorithms for admission control.

2. Online Algorithms

Motivated by the considerations above, we first review and adapt, well-known and recent algorithms from the online literature to the admission control problem in SDN. We then test their performance under different traffic conditions to understand and highlight the strengths and weaknesses of each of them. Traditionally, online algorithms for admission control fall into two main categories: i) worst-case and ii) average-case. Worst-case algorithms are characterized by max-min performance guarantees under specific worst-case scenarios where a malicious adversary chooses the worst possible sequence of connection requests. Due to their conservative nature, they generally underperform under more standard traffic conditions. One of the first online algorithms for admission control has been AAP [3]. The algorithm is $O(\log n)$-competitive, meaning that it cannot reject $O(\log n)$ times more requests than the offline optimal (n being the number of nodes). Buchbinder et al. proposed in [2] a framework to derive algorithms for online packing and covering problems with performance guarantees in the worst-case scenario. Such framework developed the theory behind the initial intuition of AAP. On the other hand, ii) average-case algorithms show high expected performance over random traffic conditions, but cannot guarantee good performance in specific adversarial scenarios. The Primal-Beats-Dual (PBD) algorithm has been introduced by Kesselheim et al. in [3]. It computes the optimal (fractional) solution of the relaxed Linear Program (LP), by considering all the past requests and scaling the capacity of the graph. It then attempts to route the request over a path randomly selected, with a probability proportional to the value of the computed fractional solution. PBD suffers from computation complexity. Agrawal et al. [4] have proposed a fast algorithm with multiplicative updates to solve this issue. It applies to i.i.d. and random order inputs. It implements an efficient stochastic gradient descent that we adapted for our admission control problem.

3. Admission Control with Experts

As observed in practice, worst-case and average-case online algorithms for admission control are better than the naive and greedy strategy which accepts every demand until bottlenecks appears.

However, there is no algorithm outperforming all the remaining ones under all traffic conditions. Luckily, modern SDN control platforms, which are running on top of commodity servers and are built upon cutting-edge distributed computing technologies, enable the parallel execution of different algorithms to solve a single decision problem. Such algorithmic architecture, called boosting or prediction with expert advice setting [6], is commonly used in machine learning. It executes all the algorithms in parallel and attempts to track the best one in an online fashion. The bulk of the literature on experts focuses on proving theoretical performance bounds in the basic non-reactive scenario where the action taken by the decision maker does not affect the state of the system, which is definitely not the case for admission control. We identified that the Strategic Expert meta-Algorithm (SEA) [5] applies to out reactive setting. Under some stationary conditions on the system (in our case, on the traffic load on the links), SEA is proven to perform at least as well as the oracle which steadily selects the algorithm with best average performance. We finally show that this approach achieves very good performance in practice, and is able to successfully overcome the several limitations imposed by the inherent online and hard-to-predict nature of the problem at hand. We evaluate the performance of the online algorithms under realistic conditions, using the real-life dataset captured in 2006 by Uhlig et al. [7] on the GEANT network.

Bibliography

1. B. Awerbuch, Y. Azar, and S. Plotkin, "Throughput-competitive on-line routing"', in Proc. FOCS, 1993.
2. N. Buchbinder and J. Naor, "Improved bounds for online routing and packing via a primal-dual approach", in Proc. FOCS, 2006.
3. T. Kesselheim, A. Tönnis, K. Radke, and B. Vöcking, "Primal beats dual on online packing LPs in the random-order model", in Proc. ACM STOC, 2014.
4. S. Agrawal and N. R. Devanur, "Fast algorithms for online stochastic convex programming", in Proc. ACM SODA, 2015.
5. D. P. de Farias and N. Megiddo, "How to combine expert (and novice) advice when actions impact the environment?", in Advances in Neural Information Processing Systems, 2003.
6. N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. New York, NY, USA: Cambridge University Press, 2006.
7. S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon, "Providing public intradomain traffic matrices to the research community", ACM SIGCOMM Computer Communication Rev, vol. 36, no. 1, 2006.