Model-Based Control Plane for Fast Routing in Industrial QoS Network

Jochen W. Guck*, Martin Reisslein‡, Wolfgang Kellerer*
*Lehrstuhl für Kommunikationsnetze; Technische Universität München; Munich, 80290, Germany
{guck, wolfgang.kellerer}@tum.de
‡Electrical, Computer, and Energy Eng.; Arizona State University; Tempe, AZ 85287-57060, USA
reisslein@asu.edu

***Abstract*—Industrial networks demand centrally controlled Quality of Service (QoS), often in the form of hard real-time guarantees. We propose a novel SDN-based QoS control paradigm that (i) maintains an accurate network model through network calculus to avoid a control loop over forwarding and control planes, (ii) routes flows over a network of "queue links", whereby each physical network link houses multiple queue links (with different QoS levels), and (iii) manages QoS through delay-constrained least-cost (DCLC) routing on the network of queue links.**

I. Introduction

The core contribution of this poster is a novel model-based control plane for achieving hard real-time QoS in an SDN network. We introduce a novel "queue link" network topology model that represents both the underlying physical link topology as well as the priority queue structure at each link. The priority queues represent different levels of forwarding QoS at the link. A standard reactive QoS routing algorithm, such as delay-constrained least-cost (DCLC) routing [1], can then find routes through the queue link network that meet the hard real-time QoS. The route selection, thus determines both the path that an admitted flow takes through the physical links as well as the QoS queues that the flow traverses in the individual physical links.

The proposed approach is fast by avoiding a QoS control loop over the forwarding and control planes. Rather, we close the QoS control loop only over the model, state, and routing blocks in the control plane. We compare this novel queue link model-based SDN-based QoS approach through simulations with a mixed integer programming (MIP) based approach. We find that the novel model-based approach achieves nearly the same connection carrying capacity of the network and same bandwidth utilization as the computationally prohibitive MIP approach.

A. Related Work

Existing adaptive QoS control architectures involve a control loop that is closed over the forwarding and control planes. Closed-loop control over both the forwarding and control planes tends to introduce delays and measurement errors that are difficult to correct to achieve fast, accurate QoS control. Software Defined Networking (SDN) emerged recently as a flexible control plane paradigm and several studies have begun to explore the implications of SDN for QoS networking. For instance, Jain et al. [2] report on an SDN-based traffic engineering system for efficient bandwidth usage in a backbone network. Several studies, e.g., [3]–[5], have explored abstraction concepts and software components towards QoS provisioning with SDN. Policy Cop [6] is a closed-control loop system (over forwarding and control planes) based on measurements in the forwarding plane for providing non-deterministic aggregated and per-flow QoS guarantees. A similar measurement-based approach with a focus on QoS for cloud communication has been examined in [7].

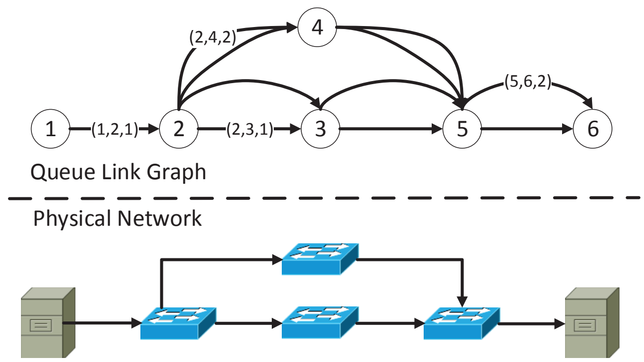

Fig. 1. Illustrative example of queue link network model: One origin node, four switching nodes, and one destination node are interconnected by physical links in the feedforward (left to right) direction. The nodes are indexed by integers 1 through 6, i.e., the set of networks nodes is $\mathcal{N} = \{1, 2, \ldots, 6\}$. Each switching node contains $Q = 2$ priority queues while source nodes have one queue. A queue link (edge) (u, v, q) connects priority queue q in node u to node v. Thus, the set of edges is $\mathcal{E} = \{(1, 2, 1), (2, 3, 1), (2, 3, 2), (2, 4, 1), (2, 4, 2), \ldots, (5, 6, 2)\}$.

II. SDN QoS with Control Plane Model

An SDN controller for industrial QoS networking poses the general problem of path computation (routing) to meet the hard real-time QoS constraints. The proof-of-concept study [8] derived network calculus foundations for evaluating the real-time QoS requirements and solved this general problem with a mixed integer program (MIP). The MIP results in prohibitive computation times already for small networks. The monolithic MIP solution approach cannot be employed for run-time decisions in realistic-sized industrial networks. In order to reduce the computational effort for solving the routing problem subject to the QoS constraints so as to operate SDN-based industrial QoS networks at run-time, we introduce a network-calculus based control plane model.

A. Queue Link Network (Control Plane) Model

In our real-time QoS framework, each switching node may be viewed as a collection of queues, each with a specific

978-1-4673-7113-1/15/$31.00 ©2015 IEEE

transmission bit rate and a specific buffer space. We define the network formed by the queues in the switching nodes and the links interconnecting the switching nodes as a *queue link network*. In particular, we denote Q_u for the number of priority queues in a given switching node u; origin nodes are considered to have one priority queue. We define the directed *queue link* (u, v, q) to connect queue q in node u via the physical link between nodes u and v to node v. As illustrated for the example of a simple feedforward network structure with two queues in each switching node in Fig. 1, the number of queue links between nodes u and v is equivalent to the number Q_u of priority queues in node u. We define the set $\mathcal{E} = \{(u, v, q)\},\ u, v \in \mathcal{N},\ q = 1, \ldots, Q_u$ to denote the set of all queue links (edges) in a given queue link network.

B. Routing

Our routing problem for a flow f corresponds to a DCLC routing problem [1], which is typically formulated to minimize the cost $C(P_i)$ of a path from origin o to destination d: $\min_{P_i \in \mathcal{P}'(o,d)} C(P_i)$ subject to the path delay $D(P_i)$ meeting the constraint t_f, i.e., $\mathcal{P}'(o, d) = \{P_i \in \mathcal{P}(o, d) \mid D(P_i) \leq t_f\}$. In our novel approach we add constraints that check the mean data rate and buffer space constraints. We define indicator functions $R(P_i, f)$ and $B(P_i, f)$ that give a one when the rate and buffer space constraints are met (and give zero otherwise). The resulting constraint is:

$$\mathcal{P}^*(o, d) = \{P_i \in \mathcal{P}(o, d) \mid D(P_i) \leq t_f \ \wedge\ R(P_i, t_f) = 1 \ \wedge\ B(P_i, t_f) = 1\}. \quad (1)$$

Based on existing DCLC routing algorithms, we developed an online algorithm for admission control, including the routing (if admissible), of a new flow f requesting admission to the network. We introduce a simple per-hop cost function (scf) and a "buffer-aware" cost function (bcf) for minimizing the end-to-end path cost. We denote P_f for the path found by a solution approach to the DCLC problem for flow f. Specifically, we denote

$$P_f = \{(u_p, v_p, q_p)\}_{p=1,\ldots,p_{\max,f}} \quad (2)$$

for the ordered sequence of queue links that constitute the path.

III. Evaluation

For the evaluation of our online routing solution based on the control plane model, we consider a typical industrial unidirectional ring structure with six switching nodes and six source/destination nodes. Each switching node output port (link) has four output queues operating with strict priority scheduling (without preemption). Each flow belongs to one of the four service classes from which we form an ensemble of 968 representative traffic mixes. For the evaluation of our online admission control algorithm, including the routing of the admitted flows, we compare with the results generated by the offline MIP in [8]. Comparisons of the computation time found that the actual online routing and admission control decision in our model-based approach requires only a few microseconds compared to the hundreds of seconds or more required by the MIP.

We show a cumulative distribution function (CDF) plot (across the 968 traffic mixes) for the total number of flows, which is equivalent to the link utilization for the considered flows with the same average bit rate, in Figure 2. The results

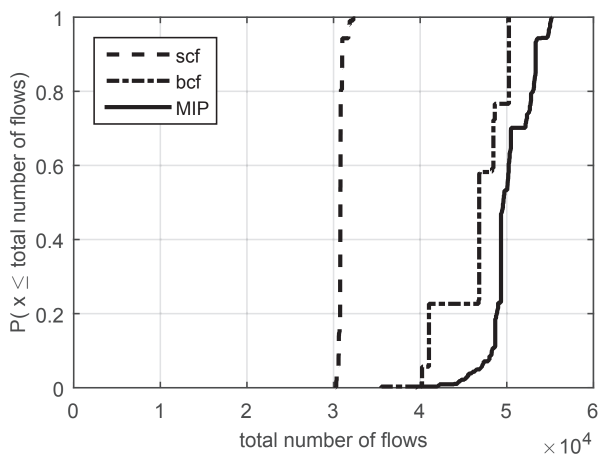

Fig. 2. Comparison of the maximal achieved total number of flows for the MIP solution and the online model-based algorithm with the simple cost function (scf) and the buffer-aware cost function (bcf).

in Figure 2 indicate higher bandwidth utilizations with the buffer-aware routing cost function. The buffer-aware routing cost function achieves up to 93 % link utilization compared to up to 58 % link utilization with the simple routing cost function (and compared up to 100 % with the MIP).

IV. Conclusion

Our novel QoS approach models the control plane so as to support online routing (and admission control) over a network of priority queues. Online delay-constrained least-cost (DCLC) routing selects the switching node as well as QoS priority queues that a flow traverses in the network. The proposed model-based approach makes accurate routing and admission control decisions within a few microseconds, compared to over hundreds of seconds required by a monolithic mixed integer program (MIP) approach. Our approach achieves up to approximately 93 % average link utilization in an industrial communication scenario compared to close to 100 % utilization by the MIP approach.

References

[1] D. S. Reeves and H. F. Salama, "A distributed algorithm for delay-constrained unicast routing," *IEEE/ACM Trans. Netw.*, vol. 8, no. 2, pp. 239–250, 2000.

[2] Jain, et al., "B4: Experience with a globally-deployed software defined WAN," *ACM Computer Commun. Rev.*, vol. 43, no. 4, pp. 3–14, 2013.

[3] I. Bueno, J. I. Aznar, E. Escalona, J. Ferrer, and J. A. Garcia-Espin, "An OpenNaaS based SDN framework for dynamic QoS control," in *Proc. IEEE SDN for Future Networks and Services (SDN4FNS)*, 2013, pp. 1–7.

[4] A. D. Ferguson, A. Guha, C. Liang, R. Fonseca, and S. Krishnamurthi, "Participatory networking: An API for application control of SDNs," in *ACM Computer Commun. Rev.*, vol. 43, no. 4, 2013, pp. 327–338.

[5] M. S. Seddiki, M. Shahbaz, S. Donovan, S. Grover, M. Park, N. Feamster, and Y.-Q. Song, "FlowQoS: QoS for the rest of us," in *Proc. ACM Workshop on Hot Topics in SDN*, Aug. 2014, pp. 207–208.

[6] M. Bari, S. R. Chowdhury, R. Ahmed, and R. Boutaba, "PolicyCop: An autonomic QoS policy enforcement framework for software defined networks," in *Proc. IEEE SDN4FNS*, Nov. 2013, pp. 1–7.

[7] A. Akella and K. Xiong, "Quality of Service (QoS)-guaranteed network resource allocation via Software Defined Networking (SDN)," in *Proc. IEEE DASC*, Aug. 2014, pp. 7–13.

[8] J. Guck and W. Kellerer, "Achieving end-to-end real-time Quality of Service with Software Defined Networking," in *Proc. IEEE Int. Conf. on Cloud Networking (CloudNet)*, Oct. 2014, pp. 70–76.