Demonstration on Fairness Among Heterogeneous TCP Variants Over 10Gbps High-speed Networks

Lin Xue*, Suman Kumar†, Cheng Cui* and Seung-Jong Park*
*Department of Computer Science, Center for Computation & Technology,
Louisiana State University, Baton Rouge, LA, USA
†Department of Computer Science, Troy University, Troy, AL, USA
Email: {xuelin, sumank, ccui, sjpark}@cct.lsu.edu

***Abstract*—Several high speed TCP variants are adopted by end users and therefore, heterogeneous congestion control has become the characteristic of newly emerging high-speed networks. In this demonstration, we show the fairness among heterogeneous TCP flows over CRON, a 10Gbps high-speed networking testbed. The significant throughput difference among heterogeneous TCP flows shows that heterogeneous TCP flows induce more unfairness than homogeneous ones. Moreover, fairness behavior changes as we change intermediate router parameters, such as queue management scheme, buffer size, etc. During the demonstration, audience will have an intuitive view on fairness behavior among heterogeneous TCP variants over 10Gbps high-speed networks.**

I. Scope of the Demonstration

High-speed networks with capacity to deliver a data transfer rate in excess of 10Gbps have been developed to serve the demand of end users. Among all the characteristics of high-speed networks, fairness among its end users has been one of the most important. A recent study [1] also shows that our traditional homogeneous network is rapidly evolving into a heterogeneous one. However, fairness among heterogeneous TCP variants has not been evaluated in a 10Gbps high-speed networking environment.

In this demonstration, we show audience fairness among heterogeneous TCP variants by using CRON [2], a 10Gbps high-speed optical network testbed.

CRON is a cyber-infrastructure with reconfigurable optical networking environment that can provide multiple networking testbeds consisting of routers, delay links, and high-end workstations operating at up to 10Gbps bandwidth. Different application developers and networking researchers can use those virtually created high speed networking and computing environments without technical knowledge of network hardware and software. With the support of NSF (award #0821741) and GENI grants, the LANET network research group at Louisiana State University has developed and operated the CRON testbed, which is based on the framework of Emulab project [3].

CRON provides integrated and automated access to a wide range of high-speed networking configurations, such as National Lambda Rail (NLR) [4], Internet2 [5], Louisiana Optical Networks Initiative (LONI) [6], as well as purely user-defined networks. Based on specific demands, users can dynamically reconfigure whole computing resources, such as operating system, middle-ware, and applications. Because of the automated and reconfigurable characteristics, all types of experiments over CRON could be repeatable and controllable.

We setup dumbbell topology with three pairs of TCP senders and receivers in CRON as shown in Figure 1. The three pairs of senders and receivers have three different high-speed TCP variants, namely TCP-SACK [7], HighSpeed TCP (HSTCP) [8], and CUBIC TCP (CUBIC) [9]. We show the significant difference of throughput among these three different TCP variants in 10Gbps high-speed networking environment. Thus, audiences can have a clear view of the fairness among heterogeneous TCP variants over 10Gbps high-speed networks.

Then, we change queue discipline in intermediate router to RED (Random Early Detection) [10] and CHOKe (CHOose and Keep for responsive flows, CHOose and Kill for unresponsive flows) [11]. We show that active queue management (AQM) schemes can tame the unfairness among three different TCP variants. Router buffer size is changed as well to observe the difference in fairness behavior. During the demonstration, we set the router buffer size to different Bandwidth-delay product (BDP), such as 10% of BDP and 20% of BDP. These different router buffer sizes show different fairness behavior to the audience.

This demonstration will showcase the configuration of 10Gbps networking environments. Fairness among heterogeneous TCP variants in 10Gbps high-speed networks will be presented to audience. More works can also be demonstrated if time allows, such as Linux parameters tuning for 10Gbps environment, queue discipline parameter setting, network delay setup, multiple long-lived TCP flows, short-lived TCP flows generation, RTT fairness among heterogeneous TCP flows, and etc.

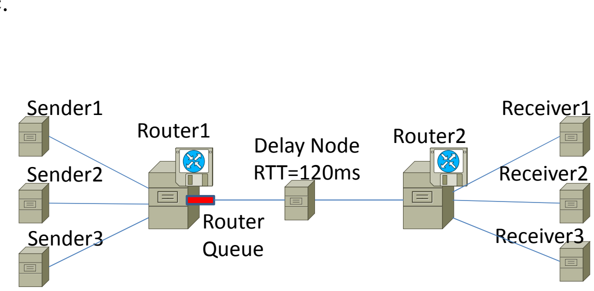

Fig. 1: Network Topology for Demo

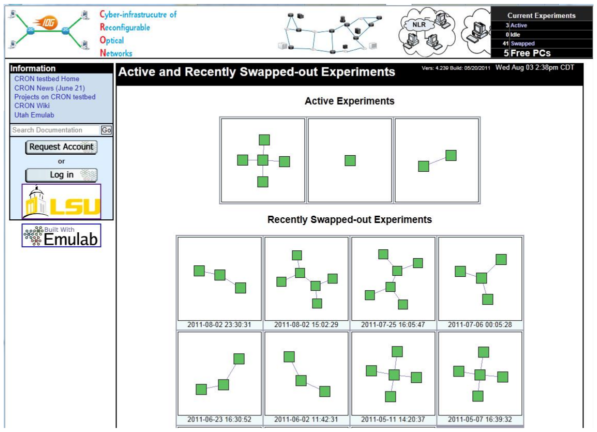

Fig. 2: Active Experiments in CRON

II. Significance of the Demonstration

To the best of our knowledge, this demonstration is the first one to show fairness among heterogeneous TCP variants in a realistic 10Gbps high-speed network environment.

Since fairness has become the most important feature among end-users in high-speed networks, it is important to develop a systematic and repeatable high-speed networking environment to evaluate and explore the fairness behavior. However, nowadays 10Gbps network environment is still scarce because of its high cost and limited access. In this demonstration, we show the experimental setup and experimental results in CRON, which is a 10Gbps high-speed network testbed. All works can be easily controlled and reproduced in 10Gbps high-speed networking environment. Network topologies and parameters are easily reconfigurable in CRON to make networking research more effective. For example, network topologies can be swapped in and swapped out within minutes; network delay can be emulated by Dummynet or hardware emulator by setting their delay parameters. A snapshot of active and recent experiments is shown in Figure. 2. Throughout the demonstration, audience can have an intuitive and straightforward experience of 10Gbps high-speed networks.

According to our findings, unfairness among heterogeneous TCP variants does exist. In comparison to homogeneous networks, the unfairness becomes more severe in heterogeneous high-speed networks. We calculate fairness among heterogeneous TCP flows in terms of the long term throughput received by each flow as Jains fairness index. To show the significant degradation of fairness while transitioning from homogeneous network to heterogeneous network, fairness index is presented in Table. I for these two cases.

TABLE I: Fairness: Homogeneous TCP vs Heterogeneous TCP (buffer size = 20% BDP, RTT = 120ms)

	Drop-tail	RED	CHOKe
CUBIC	0.988	0.994	0.991
HSTCP	0.978	0.987	0.990
TCP-SACK	0.936	0.977	0.970
Heterogeneous TCP	0.681	0.732	0.747

Also, in our demonstration, we change queue management schemes with varying degree of buffer sizes immediately in the intermediate router. We show AQM schemes improve fairness for large buffer sizes as compared to Drop-tail, whereas all three queue management schemes show similar fairness behavior for small buffer sizes. Fairness behavior among heterogeneous TCP flows combined with different router parameters is presented for the first time to the audience.

III. Basic Idea of the Demonstration

In our demonstration, firstly, we introduce the audience to a 10Gbps high-speed networking testbed, CRON. The high-speed networking environment ensures users to have their networking experiments reconfigurable, so that the experiments on fairness among heterogeneous TCP variants can be controllable and repeatable.

We use zero-copy Iperf to generate long-lived TCP streams. Three senders send three different TCP variants into a bottleneck link. Different TCP variants have different increase parameter α and decrease parameter β. In our case, TCP-SACK has α to be 1 and β to be 0.5. HSTCP uses table driven increase and decrease factors, that is α is 1 to 72 and β is 0.5 to 0.9. For CUBIC, congestion window increases based on a cubic function, which has a steadier growth. Thus, the audience will see a significant difference in throughput among different TCP variants. Fast TCP variants get more bandwidth, while slow TCP variants get less.

And then we switch our focus to the router queue. The bottleneck queue is at the output queue of Router1. We change the Linux queuing discipline in the bottleneck queue to RED and CHOKe. We show that AQM schemes tame the unfairness of TCP flows in heterogeneous high-speed networks. The reason is that AQM schemes drop packets early and gently to prevent network congestion. Thus, fast TCP flows get penalized, and slow TCP flows get protected.

We also change the router buffer size to observe the difference in fairness behavior. In case of large buffer sizes, RED

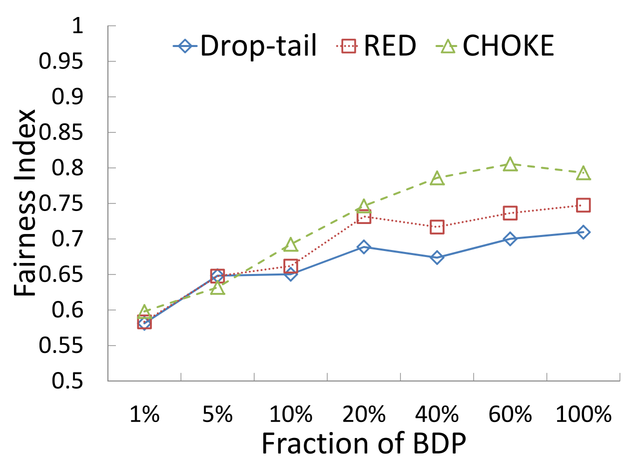

Fig. 3: Fairness for 1 TCP-SACK, 1 CUBIC, and 1 HSTCP flow (RTT = 120ms)

and CHOKe achieve de-synchronization among heterogeneous high speed TCP flows. The de-synchronization penalizes the fast flows, protects the slow flows, and improves fairness. With small buffer sizes, heterogeneous high speed TCP flows become synchronized regardless of what kind of router these TCP flows go through. As a result, small buffer size makes no differentiation among the flows, and makes all routers behave similar in fairness.

Figure 3 shows fairness index as a function of router buffer size for three different types of routers. The audience can observe the fairness behavior by throughput of long-lived TCP flow.

IV. Equipments and Setup

The setup time for this demonstration is around 10 minutes. Configuration of network topology and preparation of experiments could be done in advance of the demonstration. The whole demonstration may take 15 to 20 minutes, as each experimental showcase may take around 5 to 10 minutes.

A projector is needed to show the experimental progress and results to the audience. In addition, power and Internet/wireless access are required to showcase CRON, a 10Gbps high-speed networking testbed, as well as the experimental progress of fairness behavior of heterogeneous TCP flows.

V. Conclusion

In this demonstration, we first showcase the configuration and experimental setup on CRON, which is a 10Gbps high-speed networking testbed. Then we show fairness behavior among heterogeneous TCP variants over 10Gbps high speed networks. We also address the fairness behavior by using different queue management schemes with various queue sizes. The audience will observe that fairness problem becomes more severe in case of heterogeneous TCP flows. AQM schemes perform better in terms of fairness than Drop-tail; if buffer size in the router is small, say, less than 10% of BDP, we observe same fairness irrespective of queue management schemes. Our demonstration shows a straightforward view to audience how 10Gbps high-speed networking environment works. Moreover, we confirm that fairness among heterogeneous TCP variants depends on router parameters, such as queue management scheme and buffer size.

Acknowledgement: This work has been supported in part by the NSF MRI Grant #0821741 (CRON project), GENI grants, and DEPSCoR project N0014-08-1-0856.

References

[1] P. Yang, W. Luo, L. Xu, J. Deogun, and Y. Lu, "Tcp congestion avoidance algorithm identification," in *Distributed Computing Systems (ICDCS), 2011 31st International Conference on*. IEEE, 2011, pp. 310–321.

[2] "CRON Project: Cyberinfrastructure for Reconfigurable Optical Networking Environment," http://www.cron.loni.org/.

[3] "Emulab: Network Emulation Testbed," http://emulab.net/.

[4] "National LambdaRail," http://nlr.net/.

[5] "Internet2," http://www.internet2.edu/.

[6] "Louisiana Optical Network Initiative (LONI)," http://www.loni.org/.

[7] M. Mathis, J. Mahdavi, S. Floyd, and A. Romanow, "Tcp selective acknowledgment options," RFc 2018, October, Tech. Rep., 1996.

[8] S. Floyd, S. Ratnasamy, and S. Shenker, "Highspeed tcp for large congestion windows," IETF, RFC 3649, 2003.

[9] S. Ha, I. Rhee, and L. Xu, "Cubic: A new tcp-friendly high-speed tcp variant," *ACM SIGOPS Operating Systems Review*, vol. 42, no. 5, pp. 64–74, 2008.

[10] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *Networking, IEEE/ACM Transactions on*, vol. 1, no. 4, pp. 397–413, 1993.

[11] R. Pan, B. Prabhakar, and K. Psounis, "Choke-a stateless active queue management scheme for approximating fair bandwidth allocation," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. IEEE, 2000, pp. 942–951.