Prudentia: Measuring Congestion Control Harm on the Internet

Rukshani Athapathu, Ranysha Ware, Aditya Abraham Philip, Srinivasan Seshan, Justine Sherry
Carnegie Mellon University

1 INTRODUCTION

Growing diversity in congestion control algorithms (CCAs) raises questions about interactions between different applications in the Internet today: 1) How should we determine if a new CCA is acceptable to deploy in the Internet? and 2) Are presently deployed algorithms and services acceptable? That is, do they achieve similar performance as other connections/services sharing the same link?

To answer these questions, we utilize a new metric called "harm" [3] which quantifies the side effects or damage that a new CCA causes for competing flows and services. We present Prudentia, a new open-source, modular, and extensible system for running controlled experiments over public Internet services using the harm metric.

2 APPROACH

Traditionally, the deployability of a CCA has been viewed through the eyes of "fairness" or "TCP friendliness". However, a recent proposal by Ware et al. [3] argues that these approaches are not suitable as deployment criteria. This study argues for a harm-based approach instead: "If the harm done by a new CCA α to a widely-deployed CCA β is comparable or less than the harm done when β competes against β, we should consider it acceptable to deploy".

CCAs are always used as part of an application or service. Therefore, it is natural to suppose that CCAs are tested under the same conditions and workloads. However, much of the research related to congestion control algorithm behaviour has only considered infinite backlog flows[2, 4]. Other traffic patterns, like webpage workloads, video streaming, and teleconferencing, have largely been ignored by these studies. Unfortunately, differing application workloads and pacing can affect the fairness outcomes when two services compete for bandwidth, rendering these prior studies unsuitable to guide our understanding of practical fairness outcomes on the Internet. Hence, we propose a new testbed for measuring harm and fairness in-situ for real applications deployed on the Internet.

Now, we ask "How much harm is acceptable?" Ware et al. [3] propose using the harm caused by an already deployed, standard service and CCA - such as video streaming using TCP Reno - to quantify the acceptable harm for similar services. Since we cannot change the CCA used by a public service like YouTube, we instead develop a collection of services for video, web traffic and bulk transfers using default protocols and standard CCAs available in Linux.

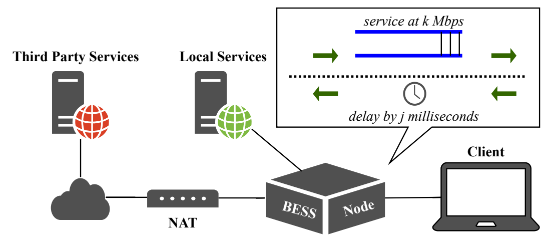

Figure 1: Prudentia testbed for local and third party services with known or unknown CCAs.

We categorize our services in three ways: the subject, victim, and baseline. The subject represents the traffic from the service we wish to evaluate for fairness/harm. The victim is the service that is harmed in the presence of the subject's traffic. Finally, the baseline represents the "default harm" that we would expect the subject to cause if it were using default protocols and a traditional CCA. Both the victim and subject can be either local or third party services. Local services represent a collection of standard applications deployed on a local server and third party services are real applications deployed on the Internet. Local services help us quantify the baseline, default harm. Once we have a baseline, we can compare it with real Internet services that have similar traffic patterns. From that we can deduce how much harm the public Internet services cause compared to the baseline harm.

3 PRUDENTIA TESTBED

Our testbed, illustrated in Figure 1, consists of three servers: a traffic sender (Third Party Services or Local Services), a traffic receiver (Client), and a node using BESS [1] software to emulate the bottleneck link. Within BESS, traffic is serviced at a configurable rate below the link capacity to introduce queuing. The bottleneck queue size is set to ratios relative to the BDP. To configure delay, we hold all ACKs for a configurable amount of time. Our testbed allows us to specify a

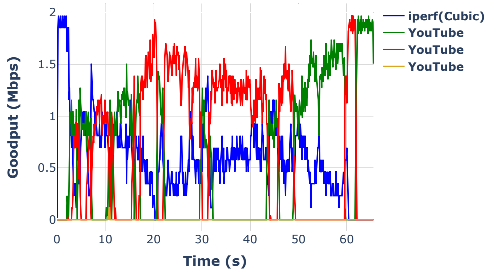

Figure 2: iperf3 (Cubic) goodput over time in competition with 1 YouTube stream which uses 3 flows on the same 20ms × 2Mbps link.

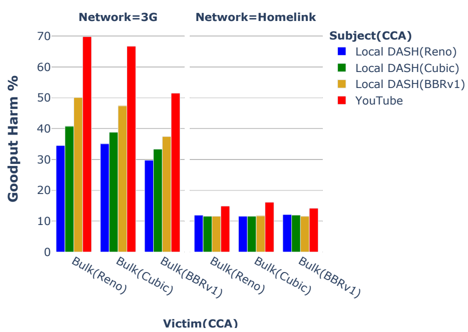

Figure 3: Goodput harm done to bulk transfers by local DASH video server and YouTube.

fixed bandwidth, fixed per-flow RTT , and a queue size to emulate different networking settings - e.g., a home broadband connection or a 3G wireless link.

We have three baseline services embedded in our testbed. To emulate bulk transfers we use `iperf3`. For webpage loads, we run a web server with a copy of the New York Times homepage (nytimes.com). For video streaming, we host a DASH video dataset on an Apache web server. We customize the Apache web server to enable configuring different CCAs per connection, so that we can have multiple web service flows competing under different CCAs simultaneously. Our testbed has the flexibility to mix and match local services with real Internet services, and can easily define the number of flows for each service as well as how long the experiment should run under which CCA.

4 CASE STUDY: VIDEO STREAMING

As a case study, we measure the harm done by a third-party video service, YouTube, to bulk data transfers. In this scenario YouTube is the subject whose harm we are evaluating, the bulk data transfer using iperf3 is the victim, and a standard DASH video service acts as the baseline (which quantifies acceptable harm). We evaluate the harm done to different CCAs by setting the algorithm used by `iperf3` and the DASH video service to Reno, Cubic or BBRv1. All experiments run for 60 seconds and repeat 10 times. We then report the median results. Our tests run under two simulated network settings. For experiments emulating a a home broadband connection, we set the bottleneck bandwidth to 50Mbps, the RTT to 20ms, and the queue capacity to 3 BDP. For experiments emulating a 3G connection, we set bottleneck bandwidth to 2Mbps, the RTT to 20ms, and the queue capacity to 3 BDP. We measure application-specific metrics for harm; in the case of bulk transfers, we measure goodput harm.

First, to show an example of how Prudentia can be used as a monitoring tool, Figure 2 presents `iperf3`'s goodput over time in competition with YouTube traffic. We can see that YouTube as a whole opens up 3 flows for data transfer, although the majority of the content is served via only 2 flows. Figure 3 shows the calculated goodput harm experienced by the `iperf3` connection in each of the aforementioned tested scenarios. From these results, we can see that YouTube does more harm to Cubic and Reno than it does to BBRv1 under both the 3G and home broadband settings. YouTube also causes more goodput harm to bulk transfers (`iperf3`) than the baseline DASH video server using any tested CCA. Hence, we can say YouTube is more harmful than a standard video streaming service would be.

5 FUTURE WORK

With the growing interest from both industry and academia to develop and deploy new CCAs, Prudentia helps in determining whether newly-developed CCAs and applications are suitable for deployment in the Internet at large. As future work, we aim to deploy an online monitoring service that will track the deployment of new CCAs, and measure their harm using our suite of tests. We suspect that this data will be of interest to watchdog groups and regulators.

REFERENCES

[1] Sangjin Han, Keon Jang, Aurojit Panda, Shoumik Palkar, Dongsu Han, and Sylvia Ratnasamy. 2015. *SoftNIC: A Software NIC to Augment Hardware*. Technical Report UCB/EECS-2015-155. EECS Department, University of California, Berkeley. http://www2.eecs.berkeley.edu/Pubs/TechRpts/2015/EECS-2015-155.html

[2] Belma Turkovic, Fernando A. Kuipers, and Steve Uhlig. 2019. Fifty Shades of Congestion Control: A Performance and Interactions Evaluation. arXiv:1903.03852. (2019).

[3] Ranysha Ware, Matthew K. Mukerjee, Srinivasan D. Seshan, and Justine Sherry. 2019. Beyond Jain's Fairness Index: Setting the Bar For The Deployment of Congestion Control Algorithms. HotNets'19, November 14-15, 2019, Princeton NJ, USA. (2019).

[4] Francis Y. Yan, Jestin Ma, Greg D. Hill, Deepti Raghavan, Riad S. Wahby, Philip Levis, and Keith Winstein. 2018. Pantheon: the training ground for Internet congestion-control research. USENIX ATC '18: Proceedings of the 2018 USENIX Conference on Usenix Annual Technical Conference. (2018).