

Comparison of TCP Congestion Control Performance over a Satellite Network

Saahil Claypool¹, Jae Chung², and Mark Claypool¹

¹ Worcester Polytechnic Institute, Worcester, MA, USA

{smclaypool, claypool}@wpi.edu

² Viasat, Marlborough, MA, USA

jaewon.chung@viasat.com

Abstract. While satellite Internet bitrates have increased, latency can still degrade TCP performance. Realistic assessment of TCP over satellites is lacking, typically done by simulation or emulation, if at all. This paper presents experiments comparing four TCP congestion control algorithms – BBR, Cubic, Hybla and PCC – on a commercial satellite network. Analysis shows similar steady state bitrates for all, but with significant differences in start-up throughputs and round-trip times caused by queuing of packets in flight. Power analysis combining throughput and latency shows during steady state, PCC is the most powerful, due to relatively high throughputs and consistent, relatively low round-trip times, while for small downloads Hybla is the most powerful, due to fast throughput ramp-ups. BBR generally fares similarly to Cubic in both cases.

1 Introduction

Satellites are an essential part of modern networking, providing ubiquitous connectivity even in times of disaster. There are 2100+ satellites in orbit, a 67% increase from 2014 to 2019 [2]. Improvements in satellite technology have increased transmission capacities more than 20x with the total capacity of planned Geosynchronous orbit satellites over 5 Tb/s.

Geosynchronous orbit satellites have about 300 milliseconds of latency to bounce a signal up and down [8], a hurdle for TCP protocols that use round-trip time communication to advance their data windows. TCP congestion control algorithms play a critical role determining throughput in the presence of network latency and loss. A better understanding of TCP congestion control algorithm performance over satellite networks is needed in order to assess challenges and opportunities that satellites have to better support TCP moving forward.

However, there are few published studies measuring network performance over actual satellite networks [17], with most studies either using just simulations [3] or emulations with satellite parameters [1, 11, 18, 19].

This paper presents results from experiments that measure the performance of TCP over a commercial satellite Internet network. We compare four TCP congestion control algorithms, chosen based on their representative approaches to

congestion control: default loss-based Cubic [15], bandwidth-delay product-based BBR [16], utility function-based PCC [11], and satellite-optimized Hybla [4]. Our network testbed and experiments are done on the Internet, but are designed to be comparable by interlacing runs of each protocol serially to minimize temporal differences and by doing 80 bulk downloads for each protocol to provide for a large sample. In addition, a custom ping application provides several days worth of round-trip time and lost packet data for a baseline satellite network with no other traffic.

Analysis of our “quiet” network gives baseline satellite loss and round-trip time characteristics. Analysis comparing the four algorithms show differences in throughput, round-trip times and retransmissions during steady state and start-up phases, with power providing a combined measure of throughput and delay.

The rest of this report is organized as follows: Section 2 presents related work, Section 3 describes our methodology, Section 4 analyzes the data, and Section 5 summarizes our conclusions and future work.

2 Related Work

Caini and Firrinielli [4] propose TCP Hybla to overcome the limitations TCP NewReno flows have when running over high-latency links (e.g., a Satellite). TCP Hybla modifies the standard congestion window increase with an extension based on the round-trip time. In Hybla slow-start, $cwnd = cwnd + 2^p - 1$ and in congestion avoidance $cwnd = cwnd + \frac{\rho^2}{cwnd}$, where $\rho = RTT/RTT_0$. RTT_0 is fixed at a “wired” round-trip time of 0.025 seconds. Hybla is available for Linux as of kernel 2.6.11 (in 2005).

Ha et al. [15] develop TCP Cubic as an incremental improvement to earlier congestion control algorithms. Cubic is less aggressive than previous algorithms in most steady-state cases, but can probe for more bandwidth quickly when needed. TCP Cubic has been the default in Linux as of kernel 2.6.19 (in 2007), Windows 10.1709 Fall Creators Update (in 2017), and Windows Server 2016 1709 update (in 2017).

Cardwell et al. [16] provide TCP Bottleneck Bandwidth and Round-trip time (BBR) as an alternative to Cubic’s (and Hybla’s) loss-based congestion control. BBR uses the maximum bandwidth and minimum round-trip time observed to set the congestion window size (up to twice the bandwidth-delay product). BBR has been deployed by Google servers since at least 2017 and is available for Linux as of kernel 4.9 (end of 2016).

Dong et al. [11] propose TCP PCC that observes performance based on small measurement “experiments”. The experiments assess throughput, loss, and round-trip times with a utility function, adopting the rate that has the best utility. PCC is not generally available for Linux, but Compira Labs³ provided us with a Linux-based implementation.

³ <https://www.compirlabs.com/>

Cao et al. [5] analyze measurement results of BBR and Cubic over a range of different network conditions, showing that the relative difference between the bottleneck buffer size and bandwidth-delay product dictates when BBR performs well. Our work extends this work by providing evaluation of Cubic and BBR in a satellite configuration, with round-trip times significantly beyond those tested by Cao et al.

Obata et al. [17] evaluate TCP performance over actual (not emulated, as is typical) satellite networks. They compare a satellite-oriented TCP congestion control algorithm (STAR) with NewReno and Hybla. Experiments with the Wideband InterNetworking Engineering test and Demonstration Satellite (WINDS) network show throughputs around 26 Mb/s and round-trip times around 860 milliseconds. Both TCP STAR and TCP Hybla have better throughputs over the satellite link than TCP NewReno – we evaluate TCP Hybla, but there is no public Linux implementation of TCP STAR available.

Wang et al. [19] provide preliminary performance evaluation of QUIC with BBR on an emulated a satellite network (capacities 1 Mb/s and 10 Mb/s, RTTs 200, 400 and 1000 ms, and packet loss up to 20%). Their results confirm QUIC with BBR has throughput improvements compared with TCP Cubic for their emulated satellite network.

Utsumi et al. [18] develop an analytic model for TCP Hybla for steady state throughput and round-trip time over satellite links. They verify the accuracy of their model with simulated and emulated satellite links (capacity 8 Mb/s, RTT 550 ms, and packet loss rates up to 2%). Their analysis shows substantial improvements to throughput over that of TCP Reno for loss rates above 0.0001%

Our work extends the above with comparative performance for four TCP congestion control algorithms on an actual, commercial satellite network.

3 Methodology

We setup a testbed, measure network baseline loss and round-trip times, serially bulk-download data using each algorithm, and analyze the results.

3.1 Testbed

We setup a Viasat satellite Internet link so as to represent a client with a “last mile” satellite connection. Our servers are configured to allow for repeated tests and comparative performance by consecutive serial runs with all conditions the same, except for the change in TCP congestion control algorithm.

Our testbed is depicted in Figure 1. The client is a Linux PC with an Intel i7-1065G7 CPU @ 1.30GHz and 32 GB RAM. There are four servers, each with a different TCP congestion control algorithm: BBR, Cubic, Hybla and PCC. Each server has an Intel Ken E312xx CPU @ 2.5 GHz and 32 GB RAM. The servers and client all run Ubuntu 18.04.4 LTS, Linux kernel version 4.15.0.

The servers connect to our University LAN via Gb/s Ethernet. The campus network is connected to the Internet via several 10 Gb/s links, all throttled to 1




Fig. 1. Satellite measurement testbed.

Gb/s. Wireshark captures all packet header data on each server and the client. The client connects to a Viasat satellite terminal (with a modem and router) via a Gb/s Ethernet connection. The client’s downstream Viasat service plan provides a peak data rate of 144 Mb/s.

The terminal communicates through a Ka-band outdoor antenna (RF amplifier, up/down converter, reflector and feed) through the Viasat 2 satellite⁴ to the larger Ka-band gateway antenna. The terminal supports adaptive coding and modulation using 16-APK, 8 PSK, and QPSK (forward) at 10 to 52 MSym/s and 8PSK, QPSK and BPSK (return) at 0.625 to 20 MSym/s.

The Viasat gateway performs per-client queue management, where the queue can grow up to 36 MBytes, allowing a maximum queuing delay of about 2 seconds at the peak data rate. Queue lengths are controlled at the gateway by Active Queue Management (AQM) that randomly drops 25% of incoming packets when the queue is over a half of the limit (i.e., 18 MBytes).

The performance enhancing proxy (PEP) that Viasat deploys by default is disabled for all experiments in order to assess congestion control performance independent of the PEP implementation, and to represent cases where a PEP could not be used (e.g., for encrypted flows).

3.2 Baseline

For the network baseline, we run UDP Ping⁵ from a server to the client continuously for 1 week. This sends one 20-byte UDP packet every 200 milliseconds (5 packets/s) from the server to the client and back, recording the round-trip time for each packet returned and the number of packets lost. Doing round-trip time measurements via UDP avoids any special treatments routers may have for ICMP packets.

3.3 Downloads

We compare the performance of four congestion control algorithms, chosen as representatives of different congestion control approaches: loss-based Cubic, bandwidth-

⁴ <https://en.wikipedia.org/wiki/ViaSat-2>

⁵ <http://perform.wpi.edu/downloads/#udp>

delay product-based BBR (version 1), satellite-optimized loss-based Hybla and utility function-based PCC. The four servers are configured to provide for bulk-downloads via `iperf3`⁶ (v3.3.1), each server hosting one of our four congestion control algorithms. Cubic, BBR and Hybla are used without further configuration. PCC is configured to use the Vivace-Latency utility function [12], with throughput, loss, and round-trip time coefficients set to 1, 10, and 2, respectively.

For all hosts, the default TCP buffer settings are changed on both the server and client – setting `tcp_mem`, `tcp_wmem` and `tcp_rmem` to 60 MBytes – so that flows are not flow-controlled and instead are governed by TCP’s congestion window.

The client initiates a connection to one server via `iperf`, downloading 1 GByte, then immediately proceeding to the next server. After cycling through each server, the client pauses for 1 minute. The process repeats a total of 80 times – thus, providing 80 network traces of a 1 GByte download for each protocol over the satellite link. Since each cycle takes about 15 minutes, the throughput tests run for about a day total. We analyze results from a weekday in July 2020.

4 Analysis

4.1 Network Baseline

We start by analyzing the network baseline loss and round-trip times, obtained on a “quiet” satellite link to our client – i.e., without any of our active bulk-downloads. Table 3 provides summary statistics.

The vast majority (99%) of round-trip times are between 560 and 625 milliseconds (median 597 ms, mean 597.5 ms, std dev 16.9 ms). However, the round-trip times have a heavy-tailed tendency, with 0.1% from 625 ms to 1500 ms and 0.001% from 1700 to 2200 ms. These high values show multi-second round-trip times can be observed on a satellite network even without any self-induced queuing. There are no visual time of day patterns to the round-trip times.

In the same time period, only 604 packets are lost, or about 0.05%. Most of these (77%) are single-packet losses, with 44 multi-packet loss events, the largest 11 packets (about 2.2 seconds). There is no apparent correlation between these losses and the round-trip times (i.e., the losses do not seem to occur during the highest round-trip times observed). Note, these loss rates are about 15x lower than the reported WINDS satellite loss of 0.7% [17].

4.2 Representative Behavior

We begin by examining the TCP congestion control performance over time for a single download representative of typical behavior for each algorithm for our satellite connection. Figure 2 depicts the throughput, round-trip time and re-transmission rate where each value is computed per second from Wireshark traces on the server.

⁶ <https://software.es.net/iperf/>




Fig. 2. Stacked graph comparison. From top to bottom, the graphs are: throughput (Mb/s), round-trip time (milliseconds), and retransmission rate (percent). For all graphs, the x-axis is time (in seconds) since the flow started.

TCP Cubic illustrates typical exponential growth in throughput during start-up, but exits slow start relatively early, about 15 seconds in where throughput is far lower than the link capacity. Thus, it takes Cubic about 30 seconds to reach the expected steady state throughput of about 100 Mb/s. During steady state (post 45 seconds) the AQM drops enough packets to keep Cubic from persistently saturating the queue, resulting in round-trip times of about 1 second. However, several spikes in transmission rates yield corresponding spikes in round-trip times above 3 seconds and retransmission rates above 20 percent.

TCP BBR ramps up to higher throughput more quickly than Cubic, but this also causes high round-trip times and loss rates around 20 seconds as BBR oversaturates the bottleneck queue. At steady state, BBR operates at a fairly steady 140 Mb/s, with relatively low loss and round-trip times about 750 milliseconds since the $2x$ bandwidth-delay product BBR keeps in flight is below the AQM queue limit. However, there are noticeable dips in throughput every 10 seconds when BBR enters its PROBE_RTT state. In addition, there are intermittent

round-trip time spikes and retransmissions from loss which occur when BBR enters PROBE_BW and increases its transmission rate for 1 round-trip time.

TCP Hybla ramps up quickly, faster than does Cubic since it adjusts congestion window growth based on latency, causing queuing at the bottleneck, evidenced by the high early round-trip times. However, there are few retransmissions. At steady state Hybla achieves consistently high throughput, with a slight growth in the round-trip time upon reaching about 140 Mb/s. Thereupon, there is a slight upward trend to the round-trip time until the queue limit is reached accompanied by some retransmissions.

TCP PCC ramps up somewhat slower than Hybla but faster than Cubic, causing some queuing and some retransmissions, albeit fewer than BBR. At steady state, throughput and round-trip times are consistent, near the minimum round-trip time (around 600 milliseconds), and the expected maximum throughput (about 140 Mb/s). The lower round-trip times are expected since round-trip time is used by the PCC utility function.

4.3 Steady State

TCP’s overall performance includes both start-up and congestion avoidance phases – the latter we call “steady state” in this paper. We analyze steady state behavior based on the last half (in terms of bytes) of each trace.

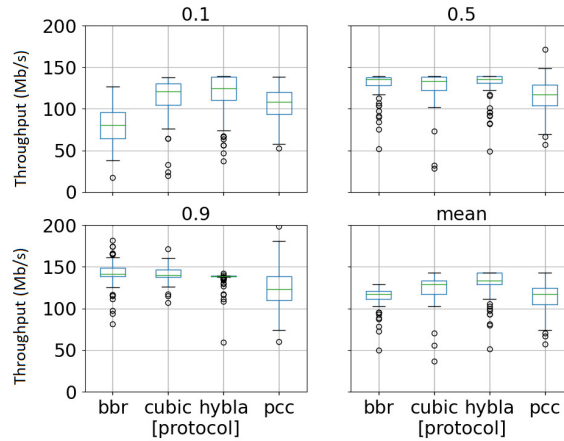


Fig. 3. Steady state throughput distributions for 10%, 50%, 90% and mean.

For each algorithm, we compute steady state throughput in 1 second intervals, extracting the 10th, 50th and 90th percentiles (and means) across all flows. Figure 3 shows the boxplot distributions. The top left is the distribution for the 10th percentile, the top right the 50th (or median), the bottom left the 90th percentile and the bottom right the mean. Each box depicts quartiles and

median for the distribution. Points higher or lower than $1.4 \times$ the inter-quartile range are outliers, depicted by the circles. The whiskers span from the minimum to maximum non-outlier. Table 1 shows the corresponding summary statistics.

Table 1. Steady state throughput summary statistics.

Algorithm	Mean (Mb/s)	Std Dev
BBR	112.9	12.2
Cubic	123.3	17.0
Hybla	130.1	17.2
PCC	112.6	17.9

Table 2. Steady state throughput effect size (versus Cubic).

	t(158)	p	Effect Size
BBR	4.44	<.0001	0.7
Hybla	2.51	0.0129	0.4
PCC	3.88	0.0002	0.6

From the graphs, at the 10th percentile BBR has lowest distribution of steady state throughput. This is attributed to its reduced throughput during the round-trip time probing phase, which, if there is no change to the minimum round-trip time, triggers every 10 seconds and lasts for about 1 second. PCC’s throughput at the 10th percentile is also a bit lower than Cubic’s or Hybla’s, possibly because PCC’s reward for a low round-trip time can result in occasional under-utilization.

BBR, Cubic and Hybla all have similar median steady state throughputs, while PCC’s is a bit lower.

BBR has the highest distribution of throughput at the 90th percentile, followed by Cubic, Hybla and PCC. BBR’s estimation of the link bandwidth may yield more intervals of high throughput than the other algorithms. Hybla’s 90th percentile distribution is the most consistent (as seen by the small box), while PCC’s is the least, maybe due to fuller queues and emptier queues, respectively (see Table 4).

From the table, Hybla has the highest mean steady state throughput, followed by CUBIC, and then BBR and PCC are about the same. BBR steady state throughput varies the least, probably since the consistent link quality provides for a steady delivery rate and round-trip time.

Since Cubic is the default TCP congestion control algorithm for Linux and Windows servers, we compare the mean throughput for an alternate algorithm choice – BBR, Hybla or PCC – to the mean for Cubic by independent, 2-tailed t tests ($\alpha = 0.05$) with a Bonferroni correction and compute the effect sizes. An effect size provides a measure of the magnitude of difference – in our case, the difference of the means for two algorithms. In short, effect size quantifies how much the difference in congestion control algorithm matters. The Cohen’s d effect size assesses the differences in means in relation to the pooled standard deviation. Generally small effect sizes are anything under 0.2, medium is 0.2 to 0.5, large 0.5 to 0.8, and very large above 0.8. The t test and effect size results are shown in Table 2. Statistical significance is highlighted in bold.

From the table, the mean steady state throughput differences compared to Cubic are all statistically significant. BBR and PCC have lower steady state

throughputs than Cubic with large effect sizes. Hybla has a higher throughput than Cubic with a moderate effect size.

Figure 4 shows the round-trip times during steady state. The x-axis is the round-trip time in seconds computed from the TCP acknowledgments in the Wireshark traces, and the y-axis is the cumulative distribution. There is one trendline for each algorithm. Table 4 shows the summary statistics.

Table 3. Baseline round-trip time summary statistics.

mean	597.5 ms
std dev	16.9 ms
median	597 ms
min	564 ms
max	2174 ms

Table 4. Steady state round-trip time summary statistics.

Algorithm	Mean (ms)	Std Dev
BBR	780	125.1
Cubic	821	206.4
Hybla	958	142.1
PCC	685	73.1

During steady state, Hybla typically has round-trip times about 200 milliseconds higher than any other algorithm, likely because its aggressive congestion window growth with high round-trip time yields more queuing delay. PCC has the lowest and steadiest round-trip times, near the link minimum, likely because its utility function rewards low round-trip times. BBR and Cubic are in-between, with BBR being somewhat lower than Cubic and a bit steadier. Cubic, in particular, has a few cases with extremely high round-trip times. Across all flows, about 5% of the round trip times are 2 seconds or higher.

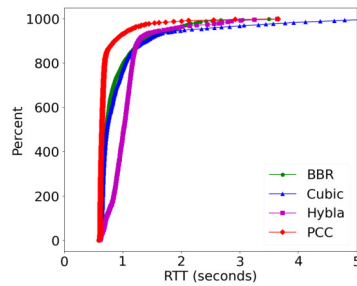


Fig. 4. Steady state round-trip time distributions.

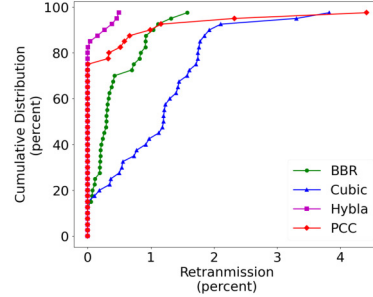


Fig. 5. Steady state retransmission distributions.

Figure 5 shows the retransmissions during steady state. The axes and data groups are as for Figure 4, but the y-axis is the percentage of retransmitted packets computed over the second half of each flow.

From the figure, Cubic has the highest retransmission distribution and Hybla the lowest. BBR and PCC are in-between, with BBR moderately higher but PCC

having a much heavier tail. Hybla and PCC are consistently low (0%) for about 75% of all runs, compared to only about 20% for BBR and Cubic.

While higher round-trip times generally mean larger router queues and more drops and retransmissions, the Viasat AQM does not drop packets until the queue is above about 1 second of delay. This means if a flow’s round-trip times remain under about 1.6 seconds, it can avoid retransmissions.

4.4 Start-Up

We compare the start-up behavior for each algorithm by analyzing the first 30 seconds of each trace, approximately long enough to download 50 MBytes on our satellite link. This is indicative of algorithm performance for some short-lived flows and is about when we observed throughput growth over time “flattening” for most flows.

The average Web page size for the top 1000 sites was around 2 MBytes as of 2018 [10], including HTML payloads and all linked resources (e.g., CSS files and images). The Web page size distribution’s 95th percentile was about 6 MBytes and the maximum was about 29 MBytes. Today’s average total Web page size is probably about 5 MBytes [13], dominated by images and video.

Many TCP flows stream video content and these may be capped by the video rate, which itself depends upon the video encoding. However, assuming videos are downloaded completely, about 90% of YouTube videos are less than 30 MBytes [7].

Figure 6 depicts the time on the y-axis (in seconds) to download an object for the given size on the x-axis (in MBytes). The object size increment is 1 MByte. Each point is the average time required by a algorithm to download an object of the indicated size, shown with a 95% confidence interval.

From the figure, for the smallest objects (1 MByte), Hybla and PCC download the fastest, about 4 seconds, owing to the larger initial congestion windows they both have (2.5x to 5x larger than either BBR or Cubic). In general, this larger initial window means Hybla downloads small objects fastest followed by PCC up to about 20 MBytes, then BBR and Cubic. After 20 MBytes, BBR downloads objects faster than PCC, perhaps because BBR exits its starting phase later than does PCC – BBR exits when the delivery rate has not increased by 25% for 3 round-trip times and PCC exits when its utility function decreases. For an average Web page download (5 MBytes), Hybla takes an average of about 4 seconds, PCC 7 seconds, BBR 10 seconds and Cubic 13 seconds. For 90% of all videos and the largest Web pages (30 MBytes), Hybla takes about 8 seconds, BBR and PCC about twice that and Cubic about thrice.

Table 5 presents the summary statistics for the first 30 seconds of each flow for each algorithm. During start-up, Cubic has a low round-trip time, mostly because it takes a long time to ramp up throughput. BBR has the highest round-trip time despite not having the highest throughput – that is had by Hybla, despite having a lower round-trip time than BBR. The relatively higher average round-trip time for BBR may be because it keeps up to a bandwidth-delay product of packets in queue. PCC has average throughputs and round-trip times, but the steadiest

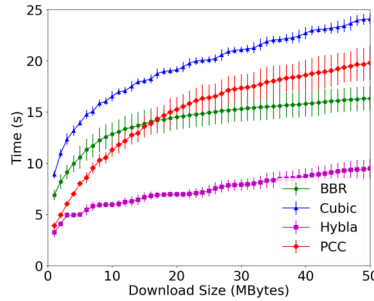


Fig. 6. Download time versus download object size.

round-trip times, possibly stabilized by the utility function rather than probing for increased data rates (and causing variable amounts of queuing) as do the other algorithms.

Table 5. Start-up summary statistics.

Algorithm	Tput (Mb/s)		RTT (ms)	
	Mean	Std Dev	Mean	Std Dev
BBR	23.1	1.8	917	42.9
Cubic	16.6	0.3	757	22.3
Hybla	40.8	2.9	799	130.8
PCC	20.3	1.6	806	15.1

Table 6. Startup throughput effect size (versus Cubic).

	t(158)	p	Effect Size
BBR	31.9	<.0001	5
Hybla	74.2	<.0001	12
PCC	20.3	<.0001	3.2

Table 6 is like Table 2, but for start-up (the first 30 seconds). From the table, the start-up throughput differences compared to Cubic are all statistically significant. The effect sizes for comparing Cubic throughput to PCC, BBR and Hybla throughputs are all very large.

4.5 Power

In addition to examining throughput and round-trip time separately, it has been suggested that throughput and delay can be combined into a single “power” metric by dividing throughput by delay [14] – the idea is that the utility of higher throughput is offset by higher delay and vice-versa. Doing power analysis using the mean throughput (in Mb/s) and delay (in seconds) for each algorithm for start-up and steady state yields the numbers in Table 7 (units are MBits). The algorithm with the most power in each phase is indicated in bold.

During steady state, PCC is the most powerful based on high throughput with the lowest round-trip times. Cubic is more powerful than BBR or Hybla since it has good throughput and round-trip times, whereas BBR is deficient in throughput and Hybla in round-trip times.

At start-up, Hybla has the most power by far, primarily due to its high throughput. BBR, Cubic and PCC are similar at about half the power of Hybla.

Table 7. TCP Power – throughput \div delay

Algorithm	Power (Mbits)	
	Steady	Start-up
BBR	145	25
Cubic	150	22
Hybla	136	51
PCC	164	25

5 Conclusion

Satellite Internet connections are important for providing reliable connectivity, but to date, there are few published research papers detailing TCP congestion control performance over actual satellite networks.

This paper presents results from experiments on a commercial satellite network, comparing four TCP congestion control algorithms – the two dominant algorithms, Cubic and BBR, a commercial implementation of PCC, and the satellite-tuned Hybla. These algorithms have different approaches to congestion control: loss-based (Cubic), bandwidth estimation-based (BBR), utility function-based (PCC), and satellite-optimized (Hybla). Results from 80 downloads for each protocol, interlaced so as to minimize temporal differences, provide for steady state and start-up performance. Baseline satellite network results are obtained by long-term round-trip analysis in the absence of our other traffic.

Overall, the production satellite link has consistent baseline round-trip times near the theoretical minimum (about 600 milliseconds) and very low (about 0.2%) loss rates. For TCP downloads, during steady state, the four algorithms evaluated – Cubic, BBR, Hybla and PCC – have similar median throughputs, but Hybla and Cubic have slightly higher mean throughputs owing to BBR’s bitrate reduction when probing for minimal round-trip times (probing ~ 10 seconds, each lasting for ~ 1.5 seconds). During start-up, Hybla’s higher throughputs allow it to complete small downloads (e.g., Web pages) about twice as fast as BBR (~ 5 seconds versus ~ 10), while BBR is about 50% faster (~ 10 seconds versus ~ 15 seconds) than Cubic. Hybla is able to avoid some of the high retransmission rates for Cubic and BBR, and to a lesser extent PCC, caused by saturating the bottleneck queue. However, Hybla does have consistently higher round-trip times, an artifact of continually having more packets in the bottleneck queue, while PCC has the lowest. Combining throughput and round-trip into one “power” metric shows PCC the most powerful at steady state, owing to high throughput and steady, low round-trip times, and Hybla the most powerful during start-up owing to fast throughput ramp-ups.

Future work includes evaluating settings for TCP, such as the initial congestion window, and algorithm-specific settings such as RTT_0 for Hybla. Since BBR does not always share equitably with Cubic [9], future work is to run multiple flows over the satellite link. When BBR v2 is out of alpha/preview, we plan to evaluate it, and QUIC [6], too. Other future work is to compare the algorithms with a performance enhancing proxy (PEP), designed to mitigate the high-latencies on the satellite link.

Acknowledgments Thanks to Amit Cohen, Lev Gloukhenki and Michael Schapira of Compira Labs for providing the implementation of PCC. Also, thanks to the anonymous reviewers and shepherd Srikanth Sundaresan for their thoughtful feedback on improving our paper.

References

1. Arun, V., Balakrishnan, H.: Copa: Practical Delay-Based Congestion Control for the Internet. In: Proceedings of the Applied Networking Research Workshop. Montreal, QC, Canada (Jul 2018)
2. Association, S.I.: Introduction to the Satellite Industry. Online presentation: <https://tinyurl.com/y5m7z77e> (2020)
3. Barakat, C., Chaher, N., Dabbous, W., Altman, E.: Improving TCP/IP over Geostationary Satellite Links. In: Proceedings of GLOBECOM. Rio de Janeiro, Brazil, (Dec 1999)
4. Caini, C., Firrincieli, R.: TCP Hybla: a TCP Enhancement for Heterogeneous Networks. *International Journal of Satellite Communications and Networking* **22**(5), 547–566 (Sep 2004)
5. Cao, Y., Jain, A., Sharma, K., Balasubramanian, A., Gandhi, A.: When to Use and When not to Use BBR: An Empirical Analysis and Evaluation Study. In: Proceedings of the Internet Measurement Conference (IMC). Amsterdam, NL (Oct 2019)
6. Cardwell, N., Cheng, Y., Yeganeh, S.H., Jacobson, V.: BBR Congestion Control. IETF Draft draft-cardwell-iccr-g-br-congestion-control-00 (Jul 2017)
7. Che, X., Ip, B., Lin, L.: A Survey of Current YouTube Video Characteristics. *IEEE Multimedia* **22**(2) (April - June 2015)
8. Cisco: Interface and Hardware Component Configuration Guide, Cisco IOS Release 15M&T. Cisco Systems, Inc. (2015), chapter: Rate Based Satellite Control Protocol
9. Claypool, S., Claypool, M., Chung, J., Li, F.: Sharing but not Caring - Performance of TCP BBR and TCP CUBIC at the Network Bottleneck. In: Proceedings of the 15th IARIA Advanced International Conference on Telecommunications (AICT). Nice, France (Aug 2019)
10. Data and Analysis: Webpages Are Getting Larger Every Year, and Here's Why it Matters. Solar Winds Pingdom. Online at: <https://tinyurl.com/y4pjrwh1> (November 15 2018)
11. Dong, M., Li, Q., Zarchy, D., Godfrey, P.B., Schapira, M.: PCC: Re-architecting Congestion Control for Consistent High Performance. In: Proceedings of the 12th USENIX Symposium on Networked Systems Design and Implementation (NSDI). Oakland, CA, USA (2015)
12. Dong, M., Meng, T., Zarchy, D., Arslan, E., Gilad, Y., Godfrey, B., Schapira, M.: PCC Vivace: Online-Learning Congestion Control. In: Proceedings of the 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI). Renton, WA, USA (Apr 2018)
13. Everts, T.: The Average Web Page is 3 MB. How Much Should We Care? Speed Matters Blog. Online at: <https://speedcurve.com/blog/web-performance-page-bloat/> (August 9th 2017)
14. Floyd, S.: Metrics for the Evaluation of Congestion Control Mechanisms. RFC 5166 (March 2008)

15. Ha, S., Rhee, I., Xu, L.: CUBIC: A New TCP-Friendly High-Speed TCP Variant. *ACM SIGOPS Operating Systems Review* **42**(5) (2008)
16. N. Cardwell and Y. Cheng and C. S. Gunn and S. H. Yeganeh and Van Jacobson: BBR: Congestion-based Congestion Control. *Communications of the ACM* **60**(2), 58–66 (Jan 2017)
17. Obata, H., Tamehiro, K., Ishida, K.: Experimental Evaluation of TCP-STAR for Satellite Internet over WINDS. In: *Proceedings of the International Symposium on Autonomous Decentralized Systems*. Tokyo, Japan (Mar 2011)
18. Utsumi, S., Muhammad, S., Zahir, S., Usuki, Y., Takeda, S., Shiratori, N., Katod, Y., Kimb, J.: A New Analytical Model of TCP Hybla for Satellite IP Networks. *Journal of Network and Computer Applications* **124** (Dec 2018)
19. Wang, Y., Zhao, K., Li, W., Fraire, J., Sun, Z., Fang, Y.: Performance Evaluation of QUIC with BBR in Satellite Internet. In: *Proceedings of the 6th IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. Huntsville, AL, USA (Dec 2018)