

Rebuttal to “Experimental Evaluation of Cubic-TCP” by Leith, Shorten and McCullagh


Injong Rhee
Department of Computer Science
North Carolina State University

We are writing this paper as a response to a paper appeared in PFLDnet 2007 authored by Leith et al. on the analysis of CUBIC [1] (hereafter called LSM). There are a few debatable issues that might mislead readers about the behavior of CUBIC. In general, the authors presented their observations on the performance of CUBIC under a special set of testing environments. However, as this paper is about CUBIC, a more balanced treatment on whether these observations are applicable uniquely to CUBIC would clarify potential misconceptions about CUBIC that LSM might cause – in other words, the authors did not discuss whether the observations are really indications of unique features of the protocol or the “normal” behaviors commonly observed in standard TCP and other HS protocols under the same environments. Note that there exists no absolute performance metrics that can define a protocol to be “unfair” or “slow” as these are all relative terms. Thus, their discussions alluding to the “gross unfairness” and “very slow convergence” of CUBIC without proper comparison with other protocols might lead to a misconception that CUBIC is really an incorrectly behaving protocol. Although the authors choose to omit such a balanced discussion or simply did not experiment with other protocols for comparison, a proper qualification of their statement, e.g., “so and so behaviors MAY NOT be unique to this protocol,” would prevent such a misconception. In addition, the authors seem to use CWND extensively as a metric to measure unfairness. But this metric can be misleading as CWND is not the same as transmission rate especially under large buffers. The purpose of this rebuttal is to point out the areas of the document that might mislead readers. As a disclaimer, this rebuttal is not reviewed by any third-party.

Detailed rebuttal:

1. CUBIC is the official name of the protocol, but not “Cubic-TCP”.
2. Fig. 4 of LSM shows slow convergence of CUBIC under the presence of 100 web flows. The characteristics of the web flows are important to understand the level of statistical multiplexing in the network. Just having 100 web flows (whose sizes are mostly small) is not going to create a lot of variations in flow rates according to our experience. Our experience tells that under some degree of multiplexing, CUBIC shows good convergence. Here are the results of such a test. The background traffic is generated by SURGE (93% log-normal body, 7% Pareto tail. The minimum file size of the tail is 133KB. The shape parameter of Pareto is

1.1. The mean and standard deviation of the log-normal distribution are 9.357 and 1.328, respectively. The arrival rate of flows is governed by an exponential distribution with intensity 0.2).



3. The authors argue that slow convergence affects the transfer rates of short-size flows in page 4. However, most short web flows tend to finish their transmission during their slow start. Besides, with enough statistical multiplexing, convergence of CUBIC flows is much faster.

4. LSM uses Fig. 5 to demonstrate that CUBIC flows can be unfair to each other. This figure is grossly misrepresenting the CUBIC behavior by showing only cwnd plots. Under 100% BDP, even difference in CWND does not make much impact if queues are saturated; thus a fair representation of this figure is to show the instantaneous transmission rate of individual flows. But the authors have conveniently neglected to plot transmission rates. Below we show a result of a similar test where we show both CWND and transmission rate of individual flows. From the figures, we can find that even if CWNDs can vary a lot under a large buffer (here in this test, approximately 6MB), their corresponding transmission rates do not vary as much. To see this, suppose the buffer is 100% BDP and one flow is running. Then the flow's CWND has the maximum at 2 times of the BDP. Therefore, even if CWND gets reduced by half at the maximum, the transmission rate still remains the same. In fact, the design H-TCP (the protocol designed and promoted by the authors) takes advantage of this fact, and it is hard to believe that the authors have missed this point. LSM uses this

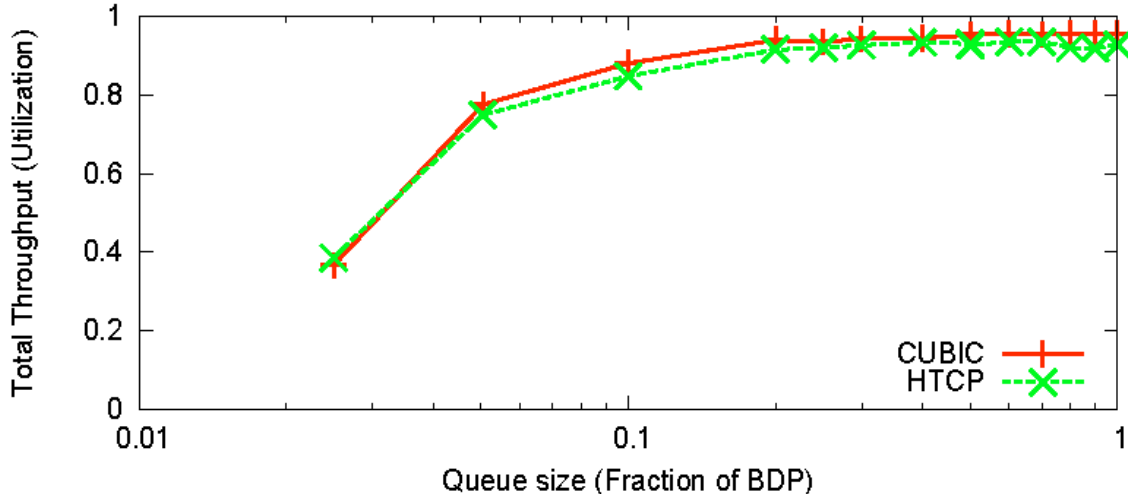
“information hiding” conveniently again in Fig. 8 where they only show CWND plot under 50 on-off sessions.





5. In page 5, LSM makes the following incorrect statement: “When a Cubic TCP flow misses a drop, the cubic function continues past the inflection point and

probes for additional bandwidth at a much a faster rate than flows recently experiencing a drop.”

6. In CUBIC, the flow, say A, that experiences a loss increases much faster than the one, say B, that misses the drop. That is because B is at its convex growing region where the beginning growth rate is very slow and A is at its concave growing region where the beginning growth rate is very fast.
7. In page 5 footnote, LSM claims that the H-TCP growth function is less aggressive than CUBIC's. This again is an optical illusion. The fair comparison of the window growth rate must be done from the point when CUBIC has the inflection point, not before. So if you move the starting position of H-TCP's to the inflection point of CUBIC's then we can see that H-TCP can be more aggressive for some period.
8. In Fig. 9, LSM shows a utilization graph where CUBIC under small buffers (in the environment without any statistical multiplexing) and a 250 Mbps link has low utilization. We have run the same tests many times and we cannot recreate the same result. We do not know how LSM generates this graph. We asked for a script for this graph to the authors, but we have not received it. Below we show our result. Below, we show our result and we also compared the CUBIC's utilization with that of H-TCP.




9. In Fig. 11, LSM shows the CWND plot of two flows with different RTTs. Again, LSM makes a good use of information hiding. In this experiment, when RTTs are different, the size of CWND does not have much meaning. The fairness needs to be measured in terms of throughput. It is unclear why the authors make this mistake over and over. The below figure shows the corresponding transmission rate graphs. Even if the two flows have different RTTs (differ by 10 times), the throughput shares are very reasonable while their CWNDs are quite different.




Again when one flow has its RTT very long, its CWND size does not make a good representation of its transmission rate. Thus, even if they could show much difference in CWNDs compared with those of short-RTT flows, their throughput difference may not be so different.


10. Fig. 12 shows two CUBIC flows to illustrate “gloss unfairness” of CUBIC. Here again, the authors conveniently choose to show throughput instead of CWND. At first, when we inspect the graph, we thought there must be some bugs in the CUBIC implementation in Linux. However, we found out that it is again a measurement error. The sampling of rate is so coarse that the application measured instantaneous rates are not shown. But if we measure the throughput from the router, we can see the actual transmission rates. In this test environment where BDP is so small (13KB), CUBIC is behaving like TCP. Therefore, CUBIC shows the same RTT fairness as TCP-Reno (with synchronization, they will have throughput ratio proportional to the square of the inverse of RTT ratio). In fact, H-TCP also has a similar behavior.



CUBIC flows (green one has RTT 160ms and red one has RTT 16 ms. The light green is total throughput measured at the router).




H-TCP (under the same environment).



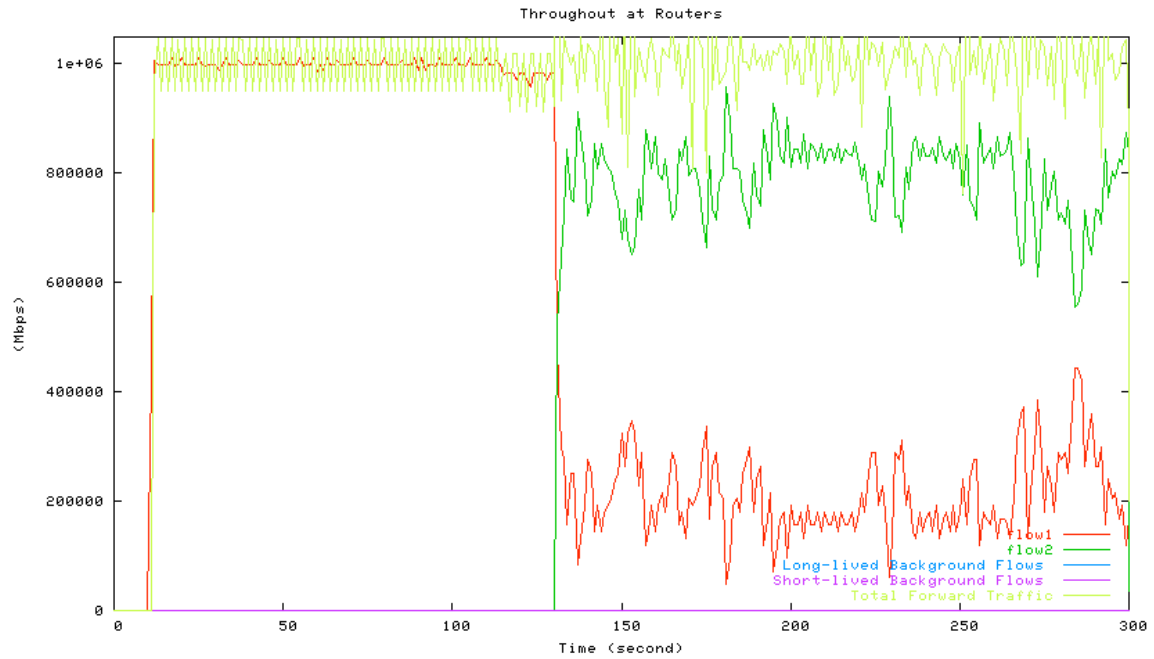
TCP-NewReno (under the same environment).

11. Fig. 14 shows the same graph as Fig. 12 but with another TCP-NewReno flow. Our result is:



CUBIC (16ms) vs. TCP-Reno flow (162ms)

With a slight increase of the router buffer to 26KB, we get the following.



In both cases, H-TCP and TCP-NewReno behave the same way.

[1] Experimental Evaluation of Cubic-TCP, D. Leith, R. Shorten, G. McCullagh, PFLDnet 2007