Consolidated Review of

A Measurement-Based Study of Multipath TCP Performance over Wireless Networks

1. Strengths

This is the first evaluation of MPTCP that I know of. The results are timely given the standardization of MPTCP. The results add a layer of realism to the previous testbed-type assessments of MPTCP.

The opportunity for multiple networks with mobile phones is currently real. The paper presents tests over live and popular network paths. The paper is sound, methodical and workmanlike. The results are well supported. Comprehensive metrics: download time, RTT; as well as various MPTCP congestion configurations. A well-rounded measurement study across various file sizes and different types of wireless networks.

The results are interesting.

2. Weaknesses

The results turn out to be basically as expected. Therefore, the paper doesn't feel fresh. Such a study might be better suited to a journal than a conference---leaving more time at the conference for fresh and new insights.

Although today there is benefit for MPTCP over 4G and Wi-Fi, is this scenario likely to continue?

The paper mentions the same observation repetitively, lots of redundancy. A short 6-page paper can easily cover the same amount of information.

The comparison across various options/configurations is not fair/right. It uses the average result across time and location, thus the variance is often very large. The results across different options largely overlap with each other. No details on the # of measurement locations and the # of measurements over time. This is critical since wireless signals vary significantly over time and location; it is unclear how 4-path MP is achieved since each laptop only has one Wi-Fi and one cellular interface.

Paper needs to specify the characteristics of the Wi-Fi networks at the different measurement locations and times (home/hotspot, a/b/g/n/ac, load). This can be important to understand why the observed losses are so high for Wi-Fi networks (it makes sense in a hotspot, but not in a home/enterprise environment). The loss rates for the Wi-Fi, whether they are consistently uniform or not, needs to be characterized.

MP-TCP has the option to use multiple interfaces in 'simultaneous mode' or 'backup mode'. This paper uses only the simultaneous mode, but due to energy considerations, the backup mode is more likely to be used. It would have been useful to look into this aspect more within the scope of this study.

While the measurements do leverage real network paths and that adds a sense of fresh sense of realism to the results, there are two drawbacks:

- ❖ The measurements are taken with one server and three static client locations and hence while the paths are real, they are limited, as well.
- ❖ The traffic is fairly simplistic. While the file size range is large (which is good), the results do not readily map to things like web transactions. I.e., web pages are large collections of objects that come from a variety of servers (origin servers, web farms, CDNs, ad networks, etc.). So, while in aggregate the pages are large, the components are often not and hence my hunch is that the performance improvements are not all that great in these cases because we're closer to the small file situations presented in the paper. However, the paper does not empirically speak to this fairly large and obvious use case.

Note: While I do think these are weaknesses, I think that in some sense they are second order and believe the strengths outweigh these weaknesses on balance.

3. Comments

This paper presents a careful evaluation of MPTCP over dual 3G or 4G and Wi-Fi networks. It's a strong paper, because it's the first careful evaluation of this protocol I've seen, and it finds a scenario where there is a real benefit.

This paper seeks to understand the performance of MPTCP for small-size flows (which is, as the authors claim, a relatively unexplored area). The result on this matter is simple: when the flow size is small, Wi-Fi often finishes download before cellular finishes initialization. When the flow size is large, mptcp uses cellular because its loss rate is much better (and that often LTE outperforms Wi-Fi in throughput). While the findings are interesting (although very intuitive), the amount of (useful) information in the paper is very limited.

This is a nice paper that helps the reader identify latency issues when dealing with single-path TCP or multi-path TCP over cellular and Wi-Fi networks. The interesting use cases arise from the variations between 3G and LTE for large file transfers. It appears that if one could make the distinction between file sizes in advance of the transfer, knowledge of the network types available, adaptively choosing the right single-path TCP over the right network can help optimize the download times, even without the use if MP-TCP, though this reviewer is conscious that such an effect cannot be guaranteed at all times.

The paper ignores one technical point that it should address: There are benefits to using Wi-Fi and 4g today, at least over Verizon. However, the benefits are minimal for AT&T and Sprint 3G--according to figure 3, for those networks, most traffic goes over one network or the other. Generalizing: my view, from your data, is MPTCP only makes sense when you have two networks that are independent AND roughly throughput equivalent, AND you have enough data to send that the setup overhead is worth the effort. Already today, this scenario happens in only one of your three cases. In the future, is it likely to be common, or more of an exception? My guess is that, in most cases, there is a clear winning network. I think the paper needs to speak to how common cases are where MPTCP will be beneficial.

The authors spent >4 pages explaining how file sizes affect the traffic partition between cellular and Wi-Fi, by first looking at a rough set of file sizes, and then a refined (and smaller) file sizes. Yet the conclusions are pretty much the same for these two sections. The RTT should be interesting but instead is a separate measurement of Wi-Fi and cellular RTT which existing works have covered. What is really interesting is whether there is any pattern that can allow MPTCP to determine whether to turn on one interface or multiple interfaces, and how to partition data between these interfaces. But this paper did not address any of these "real" problems, and just offered some straightforward measurements.

Some other suggestions:

- "NewReno", "new reno", etc. Pick one and be consistent.
- [23] has been updated and obsoleted by RFC 5681.
- 3.1: I am unclear why you set the ssthresh to 64KB. Why not infinity as recommended by the standard?
- Sec 3.1: please comment on provisioning to the WAN from your host and your wife connection point. (Having 2 1Gbps Ethernets is useless if they're plugged into a 1Mbps DSL line.) Your host at UMass is probably well provisioned, but you make no comments on the provisioning behind your remote Wi-Fi locations.
- Also: English: is "penalty" suitable for "penalization" (in "penalization mechanism"? And too much Linux in: "We also use Linux's default initial window size of Linux of 10 packets and set the default slow start threshold to 64 KB"
- Fig 2 (and other figures); you report download time on long scale, but for most larger files, what really matter is download RATE. RATE should let you switch from log scale to linear, and makes it much easier to compare different file sizes.
- I find it interesting that in section 4 there is no discussion about why the SP-Sprint results are so much worse than the others on figure 2. This seems to scream for some explanation.
- Sec 4.1 / Effect of subflow number: so most of your benefit is not actually two interfaces, but two slow starts. Web browsers have done 4 concurrent connections for this reason since 1995. Can you separate out this factor?
- 4.1: You note that the Wi-Fi in the coffee shop has "very high" traffic load. What does that mean? This is so subjective that it's basically useless. Do you have some assessment of how much traffic is going across the network? Even something crude like "X people had open laptops and Y people were seen using phones that were perhaps on the Wi-Fi" would be better than this nearly vacuous assessment.
- The plots are very difficult to read when printed. (I don't often print papers these days, but I did this one because of my own logistics.) The smaller plots like figure 5 are essentially impossible to understand in black and white printed form.
- I found 4.1.1 to be somewhat confounding. The motivation is to say that even small files should be helped by MPTCP. But, then the conclusion is that, hey! they really aren't and because of the big initial window that seems to be expected. So, it seems you're saying your own motivation was bogus. You might frame the motivation as more of a hypothesis and then show why that hypothesis doesn't hold up. Or, something. The whole \subsubsec is weird.
- Sec 4.2: Another benefit of 4G is its low loss rate (Fig 10). Can you predict what performance would be if loss rate was consistent? (The paper "Modeling TCP Throughput: A Simple Model and its Empirical Validation" from SIGCOMM 1998 might be helpful.)
- This reviewer is also wondering what was the type of Wi-Fi network observed? Was it a/b/g/n? If LTE over 3G performance is so uniquely distinct in terms of MPTCP performance, it appears that going from a slower Wi-Fi to a faster Wi-Fi standard will make a similar difference. Any insights on this will help, but specifying this in your measurement data can help the reader infer this as well.
- In 5.1 you might add a little discussion about bufferbloat and compare with the work from UNC in last year's IMC (which was directly about bloat in mobile networks). (Somewhat less similar, but still perhaps comparable is Allman's bloat assessment that is focused on non-mobile Internet traffic from CCR.)

4. Summary from PC Discussion

The paper was discussed at the TPC meeting, and it was recognized that while there were some weaknesses, the overall value of the paper was strong and justified publication at IMC. The PC does suggest that the authors address the number of weaknesses pointed out by the reviews in the camera-ready phase, and include their responses to these points in their response to the public reviews.

5. Authors' Response

We would like to thank our reviewers for their thorough and insightful comments and suggestions. In response to the reviews, we have fixed the text (typo and inconsistent terms) and incorporated them with other editorial comments. Specifically, all the figures and tables have been regenerated and the font sizes are increased to match the text font size.

One of the major comments is to characterize the WiFi network we used and to understand the cause of high WiFi loss rate. We added detailed information about the home AP we used (802.11 a/b/g) across all the measurements in residential areas, together with its wired network speed in section 3.1. The AP used in the coffee shop hotspot is 802.11n. However, its wired speed through a commercial network is unknown. Note that the high loss rates in the home network settings are mainly due to the AP we used. We did a separate set of measurements by replacing the home 802.11 a/b/g AP with a 802.11n AP, and the average loss rate is reduced from 1.6% to 0.5%, which is still much higher than its cellular LTE counterparts (<0.05%). Note that if in the future one technology might evolve to a standard of faster transmission speed (eg., WiFi from 802.11b to g to n, and to ac), the fraction of traffic carried by each path might vary, but MPTCP's goal of providing robust data transport and dynamic load balancing across different networks would remain unchanged.

Several of the reviewers were wondering if, in the future, one can make clear distinction on the file sizes to be transferred (small or large), know the types of available networks in advance, or furthermore, identify a clear winning network, is it still beneficial to use multi-path TCP? If one knows which network performs the best all the time (or in advance), he can always stay with that network. However, it could be very costly or almost impossible to decide which path is the real winner as it depends on the loss rates and RTTs over each path, as well as the file size. Most of this information is not available a priori at the client, and the loss rates and RTTs can also change over time. MPTCP, on the other hand, has been shown to be very responsive to the changes in the networks by performing load balancing across different paths/networks and can use the best path without any of this information available a priori at the clients.

In this study, we seek to understand how well MPTCP performs in the real world when the option of using all of the available paths simultaneously is given. Therefore, we tried to keep our scenarios and traffic patterns simple, in order to understand MPTCP's behaviors. Thus, we did not look at cases of MPTCP backup mode for power-saving, nor did we intend to discover a mechanism to decide when to turn off a particular interface(s). We chose AT&T LTE (rather than Verizon LTE and Sprint CDMA) for our further study in section 4 with the same reason of simplicity, as RTTs of the last two carriers exhibit high variation across different file sizes (as shown in figure 12) in our measurements. As one of the reviewers pointed out, after writing the paper, we realized that using Verizon or Sprint (rather than AT&T) for detailed study in section 4 might further emphasize the strength and importance of MPTCP.

We also include additional discussions about cellular network latency and bufferbloat in section 5.1. As for using average values in our metrics across our measurements, since network quality and condition might change over time and space, we summarize download time by boxplots with quartiles and medians. We leave other metrics with averages and standard deviations in the tables as supporting information. We also added detailed measurement information (location and rounds) for each configuration (file size/congestion controller/#path) in section 3.2.

Replies to specific suggestions:

Current Linux TCP caches routing metrics for future connections to the same destination, and use the ssthresh from previous connection as current ssthresh. As we do not cache those parameters (which might degrade performance of small flows as discussed in section 3.1), we do not have a cached ssthresh for future connections and thus require a ssthresh for fair comparisons. Furthermore, as we are using cellular networks in nearly loss-free environments, a ssthresh of infinity will lead to the case where the cellular path never leaves slow start. The congestion window of the cellular path then becomes extremely large and hence suffers severe RTT inflation which might degrade the performance of MPTCP. With a default ssthresh value, MPTCP can postpone the increase of RTTs to larger file sizes (primarily in cellular networks) since entering the congestion avoidance phase allows MPTCP congestion controller to perform joint flow control and load balancing.

The difference between a MPTCP connection of N paths and N independent concurrent TCP connections is that in the former case, we are looking at the case of fetching one object from one server over multiple paths simultaneously, where the later fetches multiple objects from multiple different servers with the same number of paths.

The reduction of download time for small file sizes demonstrates the benefits of using MPTCP. To fetch multiple objects in a webpage from different servers, another set of MPTCP connections can be established to further reduce download time.

Small flows do benefit from using MPTCP (multiple slow starts, and multiple flows). However, when the file size is really small, say 8KB or 16KB, all it requires is fewer than a dozen of packets, which can be easily transmitted through the first flow within one (or two) RTT. In this case, MPTCP behaves like single-path TCP and does not harm other TCP users. On the other hand, for very small file sizes, using MPTCP can be very useful if the initial path is very lossy, as it can recover from timeout over the initial path by retransmitting the lost packets over the second path.

Last, we have included the download rates in the paper, and a description of the background traffic load in the coffee shop environment. In our measurements, Sprint is of 3G EVDO, while other carriers are of 4G LTE. Hence, the 3G single-path TCP performance is worse than its 4G counterparts.