Summary Review Documentation for

“Network Performance of Smart Mobile Handhelds in a University Campus WiFi Network”

Authors: X. Chen, R. Jin, K. Suh, B. Wang, W. Wei

Reviewer #1

Summary: This paper measures the network performance of smart mobile handhelds (MHDs) inside a university campus WiFi network based on 2.9TB of real data. It identifies several factors that affect the network performance and compares the performance of MHDs and NHDs (non-handhelds). However, most of the observations reported are either previously known or easy to guess, hence the contributions and technique depth of this paper are quite limited.

Strengths: Given the increasing adoption of smartphones, the topic is timely. The large data set enables meaningful statistical analysis. Moreover, differentiating MHDs based on their OS types provide an interesting insight into the handhelds' performance in wifi network.

Weaknesses: Most of the results are either previously known or easy to predict. Some of the conclusions may only hold for the specific campus network where the data trace is captured, and cannot be generalized. Also, the paper fails to go beyond presenting the observations from the trace to explain or highlight the underlying technical cause or issues. This limits the contributions of the paper.

Comments to authors: This paper analyzes a large data set to study the network performance of smart mobile handhelds in a university campus wifi network. Comprehensive methods are used for the performance characterization, and a rich set of experiment results are provided.

However, part of the observations reported in this paper are either previously known or easy to guess. Basic wifi performance statistics have been available for a while in the community. More detailed session number distribution, session length distribution and delay studies can be found in previous works such as “A First Look at Traffic on Smart phones” by Falaki et.al or “Usage Patterns in an Urban WiFi Network” by Afanasyev et. Al. In addition, the factors affecting TCP performance in wireless environments are also intensively studied.

The main differences between this paper and previous studies are: 1, it is a university campus network study, 2 it categories the destination of wifi communications into several groups and studies them separately, 3 it differentiates MHDs based on their OS types. Notwithstanding, why is it important to study university campus network specifically, and why are the characteristics of campus X. Chen, R. Jin, K. Suh, B. Wang, W. Wei wifi network different from other wifi networks?

It is a well-known fact that the TCP performance depends on the network condition between the sources and destinations. Hence,

it is hard to believe that the comparative results between different destination groups presented in this paper would be valid in other wifi networks.

The differentiation between different mobile OS provides good insights into the measurement goal, but that itself is not sufficient to justify a full IMC paper.

Reviewer #2

Summary: This paper takes a fresh look at WiFi traffic of mobile devices from a university network and analyses several aspects of HTTP traffic. It splits flows into three classes and measures aspects such as servicing rate (a metric introduced by the authors), RTT, local RTT, loss rate, initial congestion window, etc. The study groups content providers in 4 big categories, with Akamai and Google being the biggest providers in terms of traffic volume.

The analysis is done separately per provider, uncovering many interesting aspects of performance. The insights can be used to optimise mobile stacks and apps, but the paper does not discuss the implications in enough depth.

Strengths: Although wide-area deployments of WiFi networks have been extensively analysed before, this paper takes a fresh, different look and brings to light some interesting findings. The paper is timely: we're witnessing standardization efforts from Google to reduce the completion times of short flows (one example being the initial cwnd change); the paper can tell us what the impact of some of these changes will be. The paper's findings could be used to optimise content delivery for mobiles.

Weaknesses: The paper is generally strong. However, there are a few questions that could be explored in more depth to make it stronger (see list below). Also, a discussion should be added to summarize the lessons from the paper and how it can affect content and application providers, etc.

Comments to authors: I liked your paper, it is well written and quite interesting. I like the fact that you look at flows in three classes and at providers in 4 classes - this makes the analysis tractable and uncovers several enlightening differences.

Your paper could be quite useful to help people optimize mobile stacks and apps except you don't bring that out enough. A discussion section that analyses the implications of your key findings would make the paper stronger.

Here is a list of questions/implications that seem relevant and have not been properly addressed in the paper.

1) Web servers are pacing traffic at the application level (at least video see the ATC paper from Matt Mathis on this) how would that affect your service rate metric? Have you taken pacing into account?

2) You mention in a few places that using a larger 1 CWND wastes resources in the mobile phone. To me that's not obvious: if the app reads the data right away, the receive buffer is very small in the kernel, and only fills up when there is a loss which is quite rare. So on average it should not make a big difference to kernel memory consumption if you have a large or small advertised receive window unless a) you have too many flows and you have to somehow reserve some memory in the case all of them get losses at the same time or b) applications are slow to read. Have you analysed these effects on iOS or Android? If not, I would be cautious in making strong claims.

3) The impact of a large initial cwnd on loss rate we see from the graphs that google does poorly for medium sized flows, but very well for small and large flows. Could this be an effect of the large initial cwnd? Can you check this from your data or otherwise?

4) You also say that some small flows have high loss rate why does that happen? Is it because of the initial cwnd? Other reasons? Answering these two questions is quite important in the context of the current debate in the IETF on standardising INITCWND=10.

5) Local rtt differences - what are the reasons for those, could you dig deeper? Other possible reasons that come to my mind are a) MHDs use Wifi power saving more, going into sleep and forcing the base station to buffer the packets longer, and b) longer flows have larger bursts of back to back packets, and these may arrive at the bottleneck link (which I assume is local Wifi) back to back and increase the average RTT.

6) From figure 13 it looks like there are two levels of caching for video/images and a single cache for text. Can you verify this somehow, do we know what google's serving architecture looks like?

7) Can you somehow figure out which is the bottleneck link for these flows in your dataset (i.e. is it in the Internet or is it the local link)? Maybe you can apply techniques similar to the packet-train to figure out what's going on.

Reviewer #3

Summary: The paper describes a measurement study of campus network to compare the performance of handheld vs non-handheld devices (such as laptops).

Strengths: The paper analyzes HTTP traffic and presents several interesting findings, including, (1) the performance to some servers such as Google is good because Google servers use a higher initial congestion window, but android phones cannot take advantage of this because it has a low receive window size. (2) there is a difference in total HTTP download time between handheld and non-handheld devices, likely due to the lower processing power of handheld devices,

Weaknesses: The analysis and the results only apply to campus WiFi networks that are likely well provisioned. This makes me think that some of the results observed may not be generalizable.

Comments to authors: My main comment is that it is not clear how well provisioned the campus wireless is and how that can

contribute to many of these results (especially the high throughput numbers)

I found it interesting that you compare the performance of MHDs and NHDs by not looking at end-to-end web latencies, but rather by looking at the latency of just the HTTP download (by carefully isolating HTTP from other client related delays). This ensures that the difference in results are not because of the difference in processor speeds between MHDs and NHDs. But this delay still doesn't take into account the processing time at the server, that may be different for NHDs and MHDs (this is likely the case if MHDs relay more on Akamai servers).

As a followup, it will be useful to characterize the differences in traffic between MHDs and NHDs. Are most users on MHDs looking at different websites and accessing different content, compared to NHDs?

Some of the conclusions you draw may not be accurate, or need to be explained better. For example, I don't understand why MHDs use better provisioned Akamai servers, while NHDs do not. This brings back to my earlier question—Is this because the traffic on the two networks are fundamentally different? Is there any way of verifying this?

It is not clear what the difference between local RTT and RTT is. The paper claims that local RTT is insignificant, but if there are wireless losses, then the last leg RTT is likely to be high.

Reviewer #4

Summary: The paper presents a detailed view of network performance over WiFi of mobile devices attached to a campus network. The main findings are that much of the traffic is from Akamai and Google servers, large congestion windows help improve performance, and Android devices start out with smaller than necessary initial congestion windows.

Strengths: Timely problem. Interesting dataset. Carefully conducted analysis.

Weaknesses: The paper is a collection of smallish findings that are all over the place. Results could be limited to campus WiFi networks.

Comments to authors: On one hand, you are studying a timely problem using an interesting dataset, the analysis appears carefully done, and the presentation is cogent. On the other hand, the paper is scattered, with results touching a range of issues. No one result stands out by itself.

Why do you analyze the older dataset in the paper, if you have equally comprehensive recent data? Given the pace at which the mobile OS updates occur, old data might not be representative any more. Actually, I am surprised that the only major difference you found in the two datasets was the fraction of Android traffic. In any case, if this paper is accepted (or even if it isn't), I'd suggest that you revise the analysis to use the recent data.

In a similar manner why did you focus on only one day's worth of data? Was processing all three days too much? There is clearly something different about Day 3, at least in terms of traffic quantity. Was it a holiday? Ideally, you wouldn't be culling data unless there is a good reason.

As I pored over the Akamai-Hartford results, I couldn't help but wonder which of your results are specific to UConn and which ones apply more broadly. For most clients, there may not be an Akamai server that close and over a high-bandwidth network.

Stepping back a bit, it seems that the overarching theme of this work is understanding what bottlenecks network performance to a mobile device today distance to the server, application patterns, or tcp parameter settings. But the paper never directly tackles this question, even though I believe you may have the data to do so. It'll be great if you could compute what the bottleneck of each transfer and what changes will bring the most improvement to user experience.

Reviewer #5

Summary: The paper measures the performance of smartphone's WiFi connection in a university campus for 3 days and compares it with that of non handheld devices.

Strengths: Comparison between MHDs and NHDs performance

Weaknesses: A very specific scenario involving university campus

Comments to authors: I liked the paper. But the results are too specific in scope (university campus), in that its not clear if they apply in other settings. The campus network may be too well provisioned, while others may not be.

At the end, it was not clear as to how to use these results. Some discussion on it will help strengthen the motivation for this study.

Other comments:

- Why would NHD flows contain larger number of requests than MHD flows?
- You mention that Akamai and google servers are well provisioned. Do you mean to imply others are not well provisioned? If so whats the basis for this observation?
- Why is the icwnd changing among flows even from the same server? What is it a function of? Some insight will help.
- You mention that Android server doesn't fully utilize the large receive window and that it wastes resources. For the number of flows active at a time, is this resource quite large? In other words, is it a big deal in that it lowers performance?

Response from the Authors

We have addressed most of the comments from the reviewers. Specifically, we have made the following major changes:

1. We have expanded Section 5.1 to explain why MHDs rely more heavily on Akamai and Google servers.
2. We have added Section 6.9 to report the findings from the data set collected in 2012, and explained the reason of presenting the results from the 2011 data set.
3. We have added Section 7 to discuss what findings are specific to UConn WiFi network and what can be applied to other networks.