A First Look at Traffic on Smartphones

by Falaki et al.

A First Look at Traffic on Smartphones

Hossein Falaki
CENS, UCLA

Dimitrios Lymberopoulos
Microsoft Research

Ratul Mahajan
Microsoft Research

Srikanth Kandula
Microsoft Research

Deborah Estrin
CENS, UCLA

Abstract—Using data from 43 users across two platforms, we present a detailed look at smartphone traffic. We find that browsing contributes over half of the traffic, while each of email, media, and maps contribute roughly 10%. We also find that the overhead of lower layer protocols is high because of small transfer sizes. For half of the transfers that use transport-level security, header bytes correspond to 40% of the total. We show that while packet loss is the main factor that limits the throughput of smartphone traffic, larger send buffers at Internet servers can improve the throughput of a quarter of the transfers. Finally, by studying the interaction between smartphone traffic and the radio power management policy, we find that the power consumption of the radio can be reduced by 35% with minimal impact on the performance of packet exchanges.

Categories and Subject Descriptors
C.2.5 [Computer-Communication Networks] Local and Wide-Area Networks – Internet

General Terms
Measurement, Performance

Keywords
Smartphone traffic, Power management

1. INTRODUCTION

Smartphone traffic represents an increasingly large share of Internet traffic. Cellular traffic is projected to grow 10 times faster than fixed Internet traffic [22] and most of this traffic is generated by smartphones [9]. By next year, smartphone sales are projected to surpass desktop PCs [18].

However, little is known today about the nature of smartphone traffic. Two recent studies have shed valuable light on some aspects of this traffic. Trestian *et al.* study the kinds of Web sites accessed at different times of the day [20]; and Maier *et al.* study HTTP traffic generated by mobile handheld devices (which include music players and personal gaming consoles in addition to smartphones) in homes [15]. Both studies are based on data gathered at a link in the middle of the network. As a result, while they can analyze traffic from a large number of devices, they do not capture a detailed, comprehensive view of individual devices. For instance, the second study misses traffic exchanged by devices through the cellular interfaces or outside of their homes.

In this paper, we report on our ongoing work on detailed characterization of smartphone traffic. Our approach is complementary to that of previous studies—we employ passive sniffers on the device and record all sent and received traffic. Given the difficulty of deploying continuous monitoring on a large number of end user devices, the breadth achievable using this method is limited. But the comprehensive view of smartphone traffic that it provides for monitored devices enables inferences that would otherwise be impossible to make. For instance, we can study how much total traffic a device generates in a day and interaction of its traffic patterns with radio power management.

The results in this paper are based on two datasets. Our primary dataset consists of 10 users across two smartphone platforms. For these users, we deployed a logger that captured packet-level traces. Our other dataset consists of 33 Android users. It contains bytes sent and received by each application in every two-minute window. We have from 1 to 5 months of data for each user.

Using these datasets, we shed light on several aspects of smartphone traffic. By analyzing commonly used ports and applications, we quantify traffic generated by various applications. We find that browsing contributes over half of the traffic, while each of messaging (email, IM), maps and media contribute roughly 10%.

We also find that most smartphone data transfers are small, with the median size being only 3 KB. Such small transfers have many implications. For instance, the overhead of lower layer protocols can be high. We show that for half of the transfers, header bytes constitute over 12% of the total bytes. In the presence of transport security, this overhead grows to 40%. For half of the transfers, lower layer handshakes constitute 20% of the total completion time.

Consistent with controlled experiments with probe traffic [12], we find that smartphone data transfers experience high delays and losses. Unlike controlled experiments, however, our data allows us to also study the impact of these path characteristics on actual smartphone workloads. Focusing on transfers with more than 10 data packets, we find that the median throughput is only 3.5 kbps in the downlink (from the network to the smartphone) and 0.8 Kbps in the uplink.

Analysis of what limits the throughput of smartphone data transfers [23] reveals that packet loss is the primary culprit. But interestingly, a quarter of the downlink transfers are bottlenecked by the size of the sender-side transport buffer. The throughput of such flows can be improved by simply increasing the buffer sizes at servers that communicate with smartphone clients.

Finally, we study the interaction of smartphone traffic with the radio power management policy. We find that the current sleep

IMC'10, November 1–3, 2010, Melbourne, Australia.

Andrew Zafft
CS Department

Agenda

- Objective
- Study Structure
- Outcomes & Observations
- Future Work / Citations
- Conclusions

Objective

Statistics

- Why Needed?
  - Building conclusions on past events (especially complex systems)
  - Understanding the current state
  - Prediction for future events

Statistics need to be rigorously applied!

“42.7% of all statistics are made up on the spot.”

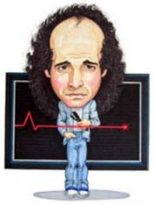

[1]

-Steven Wright

Objective

- Changes in current network traffic dynamic
  - Mobile traffic is growing 10 times faster than fixed traffic
  - Smartphones make up the majority of mobile traffic
  - Smartphone sales to surpass desktops

- The Problem
  - Prior papers studying traffic were based on a link in the middle of the network (not device level)
  - Past studies did not focus solely on smartphones

- **Solution**
  - **Capture smartphone traffic from the end-level device and analyze the data**

Study Structure

- Perform 2 independent studies on smartphone traffic
- Analyze the results looking for ways to maximize transmission efficiency
- Compare to prior network traffic studies

Study Structure

- Group 1
  - Mix of Windows Mobile & Android Users
  - Small dataset (only 10 users)
  - Packet level tracing using *Netlog* and *tcpdump*
  - Captured data sent and received down to the level of data link layer headers

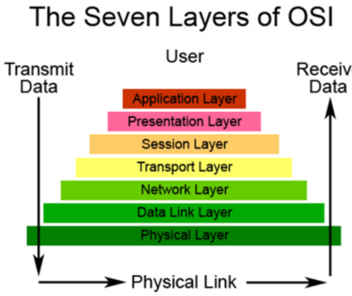

[2]

  - Captured an average of 53 days of data per user
  - Entire user group resided in 2 cities

Study Structure

- Group 2
  - – Purely uses the Android OS
  - – 33 users. While a larger user base than Group 1, this is still a fairly small set.
  - – Captured application level traffic data using a custom logging tool
  - – Recorded the number of bytes sent and received per process every 2 minutes
  - – 50 days of logging per user on average
  - – A mix of knowledge workers and high school students

Results from participants showed no valid statistical differences among the two demographics with respect to traffic, and so are reported jointly

Personal Notes

- Paper appears well designed and data appears well analyzed
- Small datasets were present, which the authors ***did*** make note of
- Overall did a good job of reviewing traffic, drawing sound conclusions and proposing better optimizations (working within their limitations)
- Authors really liked Cumulative Distribution Function

Outcomes & Observations

WiFi encourages significant bandwidth, but cellular still present

- Biases from prior work: Android users interact with their device more often than Windows Mobile users
- Traffic is roughly one order magnitude smaller than residential broadband traffic
- Dataset 2 used WiFi traffic in a much higher frequency
  - Results inferred for Dataset 1 by observing interface addresses and path delays. Dataset 2 could reliably use interface state.
- Conclusions
  - WiFi users produced significantly more bandwidth than non-WiFi users
  - Devices focused solely on cellular or WiFi bands only could miss a significant portion of the market

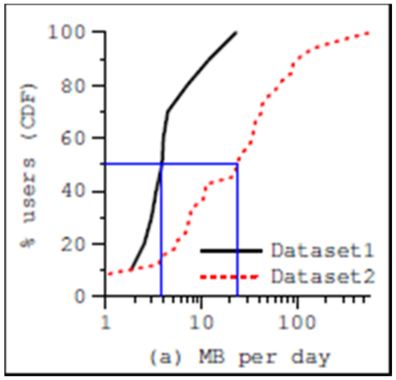

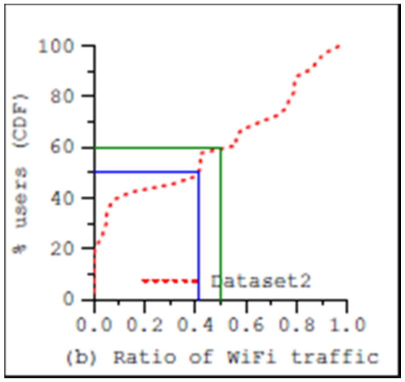

Outcomes & Observations

Browsing & email is king

Port/Application Usage

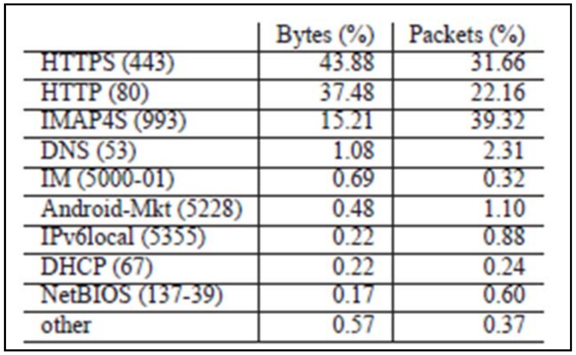

	Bytes (%)	Packets (%)
HTTPS (443)	43.88	31.66
HTTP (80)	37.48	22.16
IMAP4S (993)	15.21	39.32
DNS (53)	1.08	2.31
IM (5000-01)	0.69	0.32
Android-Mkt (5228)	0.48	1.10
IPv6local (5355)	0.22	0.88
DHCP (67)	0.22	0.24
NetBIOS (137-39)	0.17	0.60
other	0.57	0.37

Dataset 1

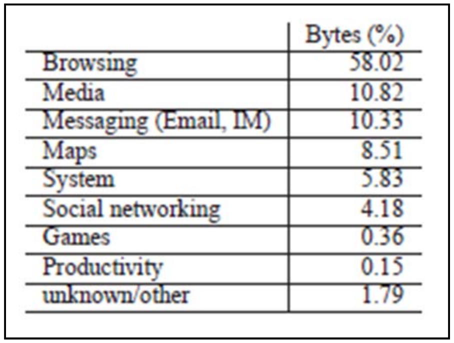

	Bytes (%)
Browsing	58.02
Media	10.82
Messaging (Email, IM)	10.33
Maps	8.51
System	5.83
Social networking	4.18
Games	0.36
Productivity	0.15
unknown/other	1.79

Dataset 2

- HTTP, HTTPS, IMAP4S and Browsing activities make up the vast majority of network traffic
- IMAP4S in particular appears to send a large number of small packets
- The preference for HTTP & HTTPS protocol could indicate the use of "tunneling" applications, resulting in misclassification of the purpose of packets

Outcomes & Observations

Download to upload consistent

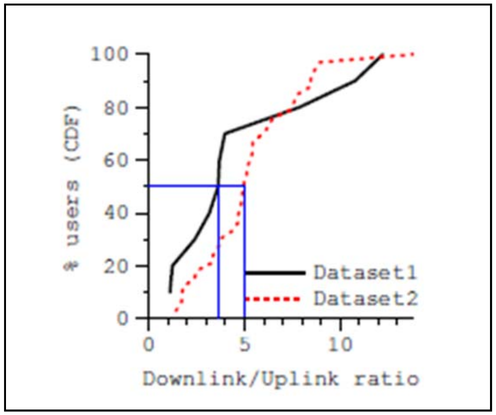

- Download to Upload ratio of 6:1
- Optimizing download activity will result in the best “bang for the buck”
- Download to upload ratios relatively similar between the two datasets.

Conclusions on Traffic Composition

- Mostly in line with similar past studies, except the composition of WiFi traffic

Outcomes & Observations

Transmissions are small

- Small mean transfer sizes: 273 KB sent, 57 KB received
- 30% of all transfers have fewer than 1 KB and 10 packets

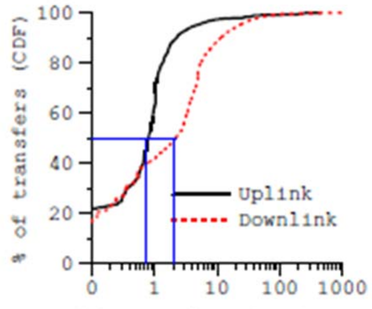

(a) Transfer size in KB

Outcomes & Observations

TCP & SSL are weighty protocols

- 96% of traffic is TCP based and more than half use SSL
- Median TCP overhead is 12%, median SSL overhead is 40%

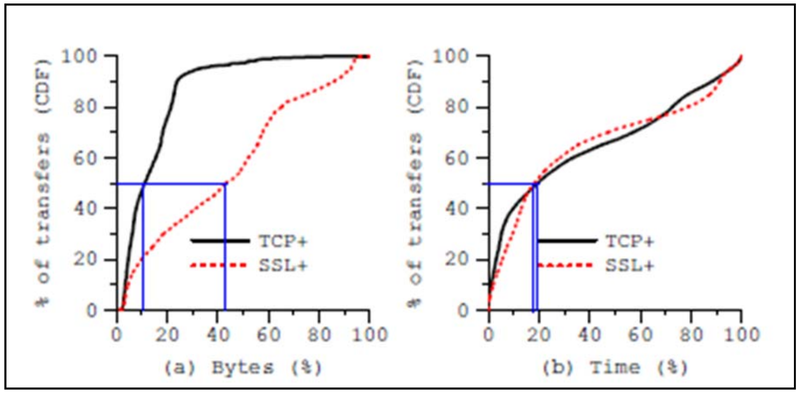

TCP | Data — Transport

IP | TCP | Data — Network

Header | IP | TCP | Data | Trailer — Data Link

Sources of Overhead [3]

- Median TCP overhead in terms of transmission time is ***20%!***
- Suggestions for future: bundle multiple transfers across applications

Outcomes & Observations

Transmission Times are Slow

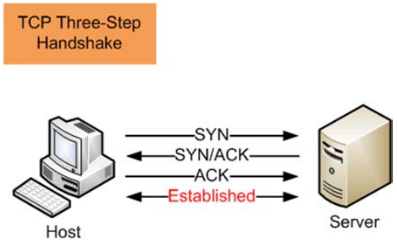
TCP Three-Step Handshake

[4]

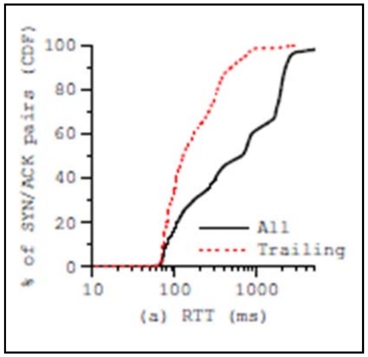

- “Trailing” corrects for radios that are asleep at beginning of transfer
- The median for trailing transfers is 125ms, with 10% of all transmissions taking over 0.5 seconds.
- When including the time to turn on radios, the median grows to 400ms and the top 10th percentile takes 1.7 seconds

- While the “trailing” time is nice to know, 1.7 seconds is what the user feels on average. For a single ACK to return in **1.7 seconds is an eternity!**

Outcomes & Observations

- Packet loss is the major culprit of delay
  - Uplink retransmission rate: 3.7%
  - Downlink retransmission rate: 3.3%

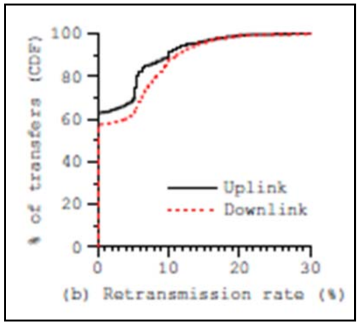

Reference Point: wired retransmission rates are less than 1%

- Roughly 40% of connections require retransmissions
- 10% of retransmissions resend 10% of their total packets

Outcomes & Observations

Throughput is low

- Median uplink rate is 0.8 Kbps
- Median downlink rate is 3.5 Kbps
- Sender window limits a quarter of download transfers

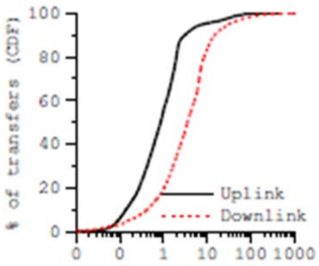

(c) Throughput (kbps)

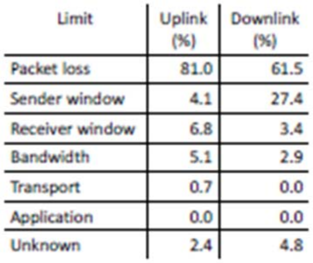

Limit	Uplink (%)	Downlink (%)
Packet loss	81.0	61.5
Sender window	4.1	27.4
Receiver window	6.8	3.4
Bandwidth	5.1	2.9
Transport	0.7	0.0
Application	0.0	0.0
Unknown	2.4	4.8

(d) Throughput limits

Conclusions: The limiting window size suggests that increasing the window on the servers will increase the download rate. These window sizes were most likely created based on wired clients (or possibly streamlined for a high volume of users)

Outcomes & Observations

Radio sleep time could be optimized

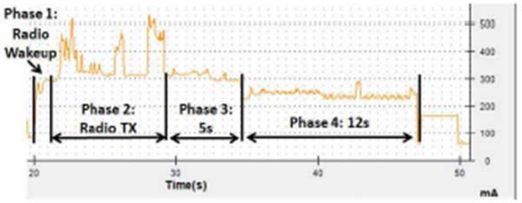

Figure 7: Current consumption (@4.2V) of HTC Touch.

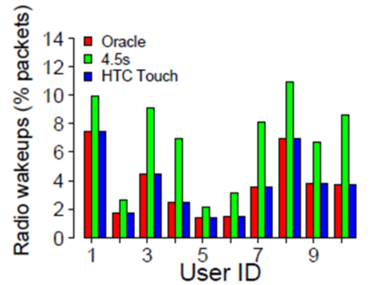

- Radios account for 1/3 the power drain on a device
- Optimal sleep time depends on burstiness of traffic
- 95% of packets are transmitted within 4.5 seconds of previous packet
- The currently implemented sleep “tail” on smartphones is 17 seconds

Conclusions: Reducing the tail to 4.5 seconds would add an additional 2-5% of packets needing to wake up the radio, but would save 35% in power consumption overall.

Future Work / Citations

- This work was completed in 2010, so no citations so far (last checked on 2-15-2011)

Conclusions to be Drawn

Smartphones suffer from many problems, all of which are sources of improvement

- High power consumption from too long sleep “tails”
- Higher than normal transmissions of small sets of data
- High overhead in transmissions
- High errors rates in transmissions

Potential Solutions

- Decrease sleep “tails”
- Group together data transmissions
- Implement better error correction procedures

Questions?

References

- [1]. http://files.myopera.com/demiphonic/blog/stevenwright50per.jpg
- [2]. http://www.windowsnetworking.com/articles_tutorials/OSI-Reference-Model-Layer1-hardware.html
- [3]. http://www.bitsontheline.net/wp-content/uploads/2009/02/encapsulation2.png
- [4]. http://humanmodem.com/images/TCP_Handshake.gif