

iPhone Usage Characteristics From Real Outdoor Wi-Fi Networks

Amalavoyal Chari (chari@tropos.com) and
Cyrus Behroozi (cyrus@tropos.com)
Tropos Networks, 555 Del Rey Ave, Sunnyvale, CA 94025

February 11, 2010

Data-hungry smartphones are highlighting the capacity constraints of commercial 2G and 3G cellular networks. However, dual-mode smartphones like the Apple iPhone are capable of switching to Wi-Fi where it is available, and this data offloading to Wi-Fi has the potential to reduce the pressure on cellular networks. In this paper we analyze data on iPhone usage on outdoor Wi-Fi networks to better characterize mobile data usage patterns and to quantify some of the offloading benefits.

1 Introduction

Residents in a handful of cities around the US have free Wi-Fi access covering their entire city. Unlike indoor Wi-Fi hot spots, which provide coverage that barely extends out to the sidewalk, a city-wide Wi-Fi mesh consists of high-powered pole-top mesh nodes covering streets and outdoor areas.

Various studies have reported that smartphones, such as the iPhone, consume an order of magnitude more data than previous generations of “feature phones”. AT&T is reported to have seen a 7000% increase in mobile data usage over the last 3 years, since the introduction of the iPhone¹. Mobile data usage has an annual growth rate of 131% and is expected to increase 66x from 2008 to 2013², as smartphone penetration and usage ramp up. The heavy mobile data usage by smartphones like the iPhone is, in turn, exposing the real capacity constraints of commercial cellular networks.

While it is difficult to build a Wi-Fi network that provides perfect coverage on its own, Wi-Fi and cellular networks can complement one another thanks to the iPhone’s ability to automatically switch between Wi-Fi and 3G. Wi-Fi provides raw data capacity while 3G fills in with more ubiquitous coverage. Since Wi-Fi networks have much higher data capacities than cellular (2G/3G) networks, they can play a useful role in offloading data traffic from the capacity constrained cellular networks. AdMob’s Mobile Metrics Report in 2008 found that 42% of iPhone requests were sent over Wi-Fi networks³.

In this study we analyze iPhone data traffic that is transported over operational Wi-Fi mesh networks covering 90+% of the outdoor areas in 11 cities around the US, and present some observations and conclusions based on the data. A large proportion of mobile data usage from dual-mode smartphones is also offloaded to indoor Wi-Fi networks in the home, in the enterprise and in hotspots, but this study focuses exclusively on data on outdoor Wi-Fi usage.

2 Networks Analyzed

We gathered data from 11 Tropos Wi-Fi networks across the country. All provide public Internet access, but three charge for access and the other eight are free. The cities are spread out across the country, including both coasts, the Midwest, and Southeast. Wealthy cities are represented, as are low-income communities. We looked at the total number of users and the numbers of unique iPhones on each network, as well as the traffic from each of these groups. Detailed location and time information was also gathered, allowing us to see usage patterns.

3 Observations and Conclusions

I. Mesh Wi-Fi networks see order-of-magnitude higher data usage than commercial cellular networks

3G cellular networks are capacity-constrained for a number of reasons, including limited licensed spectrum and a tower-based architecture that is more suited for voice than data. Even though Wi-Fi networks operate in unlicensed frequency bands, the dense-cell architecture combined with hundreds of MHz of unlicensed spectrum result in their having much higher data capacity than 2G/3G networks.

Correspondingly, the usage of real-world outdoor Wi-Fi mesh networks exceeds the usage of cellular networks by orders of magnitude. For a city the size of Mountain View, CA, the total amount of cellular data traffic, summed over all the major wireless carriers, works out to about 77 GB/day⁴. In contrast, the actual data transferred over the citywide Google WiFi network in Mountain View is 510 GB/day, which is larger by more than a factor of 6.

Another example we analyzed is the city-wide Wi-Fi network in Ponca City, OK. Free Wi-Fi access is available over the entire 18 square miles of the city, provided by 471 poletop mesh routers. Ponca City's Wi-Fi network carries ~500 GB/day (compared to a combined estimate of 30 GB/day for all the major carrier cellular networks in Ponca City). iPhones and other smartphones account for only a tiny fraction (less than 1%) of the overall data usage, with the majority coming from devices such as laptops.

Looking just at iPhone traffic, on average 530 unique iPhones connect to the Wi-Fi network each day, and consume 2.2 GB of data. Based on Ponca City's population and Apple's published sales figures, we expect there to be approximately 650 iPhones owned by Ponca City residents. This suggests that nearly every iPhone in Ponca City connects to the Wi-Fi network each day.

II. Captive portals inhibit mobile data usage

Of the free Wi-Fi networks whose usage we analyzed, four simply provide an open SSID with no additional authentication or login required, while four others require users to

click through a landing page or captive portal. As expected, the free networks carry more iPhone traffic than the paid networks, since they see a lot more users and usage. However, even among the free networks, we were surprised to see how big of an effect an open SSID with no landing page had in streamlining access for iPhones and increasing users and traffic.

We compared Ponca City (no captive portal) to another network (“City B”, to preserve anonymity) with a captive portal login. Both cities have almost exactly the same number of poletop routers and physical coverage area. Despite having nearly three times the population (and therefore three times the expected number of iPhones), City B saw a similar number of iPhones connect, and handled less than a third the iPhone data traffic of Ponca City. Most iPhone users probably didn’t tap through the captive portal after connecting, and didn’t pass substantial traffic. In other words, the presence of a captive portal seems to result in fewer users connecting and less per-user data traffic over Wi-Fi.

Although Apple’s iPhone 3.0 software release improved the user experience with regard to captive portals (bringing up a login screen from any application, not just Safari), the numbers clearly show that it doesn’t work smoothly or automatically enough to compensate for the added friction involved in interacting with a login screen on a mobile device. All else being equal, we estimate that the elimination of the captive portal on a Wi-Fi network would result in a 10-fold increase in the amount of iPhone traffic it carries (2-3x more users and 3-4x more traffic per user).

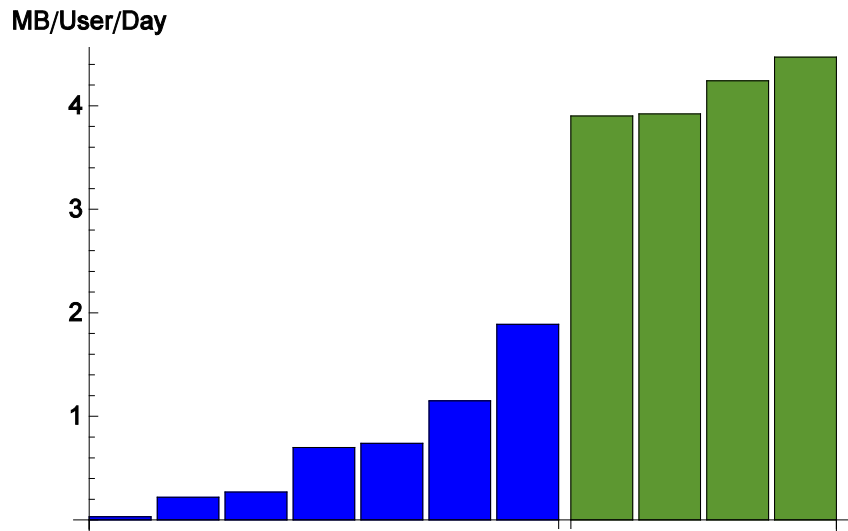



Figure 1: Data transfer rate per iPhone user (MB/day) in networks with and without captive portals

III. A small proportion of heavy users account for most of the data transfers



Previous studies of mobile data usage have noted the phenomenon of a small number of

heavy users accounting for a very substantial fraction of the data transfers. Our data provides further support for this conclusion.

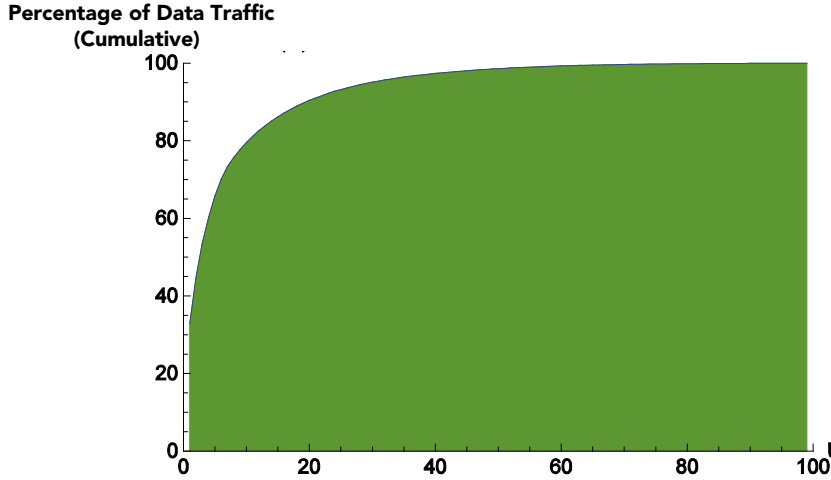


Figure 2: Cumulative distribution of iPhone user traffic

From the dataset comprising two weeks of usage on Ponca City’s Wi-Fi network, we see that the top 5% of users account for 66% of iPhone traffic, the top 10% for 80% of traffic and the top 20% for 90% of traffic. (Figure 2) The top 5% of users each transferred over 350 MB/month, the top 10% transferred over 120 MB/month and the top 20% transferred over 48 MB/month. For comparison, the Morgan Stanley Mobile Internet Report estimates that iPhones average 150-200 MB/month of data transfer, though it’s unclear whether their estimate accounts for 3G usage only or 3G + Wi-Fi.

IV. Outdoor Wi-Fi data usage tends to be concentrated in “high value” areas

While citywide Wi-Fi networks offer close to ubiquitous outdoor Wi-Fi coverage, the actual data usage by smartphone users tends to be concentrated in certain areas. We see a large proportion of outdoor iPhone connections originating from high school and college campuses, parks, playgrounds, bus stops, train stations, athletic fields, parking lots, car washes, churches, and pedestrian malls. These are all places where people use their phones, but where hot-spot access is typically unavailable.

Percentage of Data Traffic
(Cumulative)

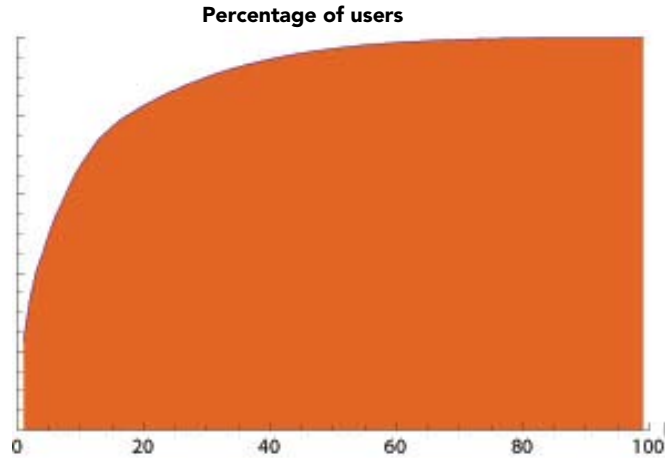


Figure 3: Cumulative distribution of iPhone traffic across routers in the network

For example, Figure 3 shows the distribution of iPhone data usage across one of the networks in our study. We see that usage is fairly concentrated within a small set of routers: the top 5% of routers (24 out of 471) account for 50% of the traffic while the top 20% of routers (or 94 out of 471) account for 80% of the data transfers. These smaller subsets of routers would comprise some of the highest value initial areas for a cellular carrier pursuing an offloading strategy to build out Wi-Fi coverage complementary to 3G. Figure 4 shows the geographical distribution of the top 5% of routers (larger circles represent the 24 highest-traffic routers and smaller dots represent the remainder of the routers).




Figure 4: Geographic distribution of the top 5% (high-traffic) routers, which account for 50% of traffic

4 Conclusion

Our measurements of iPhone data usage on outdoor Wi-Fi networks demonstrate that 3G offloading is not only feasible, but already occurring on a city-wide scale. Although all of the networks surveyed help offload traffic to some extent, the most effective are those with no captive portal and with nodes placed strategically in high value areas where significant mobile data usage occurs. We believe that Wi-Fi mesh technology is a strong complement to 3G cellular technology, and can significantly boost capacity, enhance coverage, and provide a better overall user experience.

5 Methodology

How we classified iPhones

It is fairly tricky to determine which devices on a network are iPhones, iPod touches, or other Apple devices, since the corresponding Wi-Fi MAC address ranges have not been made public by Apple. It is straightforward to determine the manufacturer of a Wi-Fi client by looking at the first three octets of its MAC address (the OUI). However, this tells us only whether or not a device is made by Apple, and not what kind of device it is. Currently, Apple has some 35 different OUIs, with hundreds of sub-ranges for different products. Often, both laptop (e.g. MacBook) and iPhone MAC address ranges are found within the same OUI. Because a laptop can consume orders of magnitude more traffic than an iPhone, a single mis-classified device could badly skew our numbers. It is therefore critical to identify devices as accurately as possible, and to err on the side of missing a few iPhones rather than accidentally including a MacBook in the iPhone data totals.

We determined the precise MAC address ranges specifically for iPhones by inspecting user-agent strings in HTTP requests and correlating these to MAC addresses. In many cases, the user-agent string contains device information identifying it as an iPhone, iPod touch, or other device. By collecting user-agent strings for over 2800 devices, we were able to deduce the precise MAC address ranges for different Apple products. The final lists contained more than 30 separate ranges each for iPhones and iPod touches—probably indicative of the number of different facilities and assembly lines that have been used to manufacture these products. In the case of adjacent ranges being used for iPhones and laptops, we chose addressing boundaries conservatively, assuming that they did not extend beyond the last observed iPhone. This ensures that our results are very unlikely to be polluted with data from laptops or other devices.

Identifying Unique Users

On any given day, a Wi-Fi network sees more iPhones connect than the number that actually pass any traffic over it. iPhones may automatically associate to the outdoor network even if the user is not engaged in any browsing activity. In addition, there are users who may encounter a captive portal or login screen and never click through past that stage of the login process, and therefore pass only a negligible amount of traffic. For our purposes, we define active users

(somewhat arbitrarily) as those that have transferred a minimum of 100 kB of traffic over the network, to avoid counting accidental associations and uncompleted connections.

Since each of the Wi-Fi networks studied covers many square miles, it is quite common for a user to connect to multiple different routers within the city over the course of the period studied. We identify each unique iPhone user by their MAC address and aggregate traffic for that user across all the routers within the city to which they connected during the 2-week period of the study. We also sum the transfers in the upload and download directions for each user to account for their overall data traffic, though the traffic in the upload direction is typically only a small percentage of the overall traffic.

References

- 1 <http://www.von.com/articles/att-connected-devices-the-ipad-and-more.html> ,
<http://gizmodo.com/5441751/okay-that-is-kind-of-a-lot-of-data>
- 2 'Cisco Visual Networking Index: Forecast and Methodology, 2008-2013',
http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html
- 3 'AdMob Mobile Metrics Report', November 2008, <http://metrics.admob.com/2008/12/november-mobile-metrics-wifi-usage-growing-rapidly/>
- 4 Morgan Stanley, 2009 Mobile Internet Report, December 2009
(http://www.morganstanley.com/institutional/techresearch/pdfs/mobile_internet_report.pdf)
estimates that the total data usage of all carrier networks in North America is currently 200,000 TB/year. We assume data traffic is proportional to population served and estimate the total traffic corresponding to Mountain View (population 70,000) to be about 77 GB/day