

A Comparative Study of Handheld and Non-Handheld Traffic in Campus Wi-Fi Networks

Aaron Gember, Ashok Anand, and Aditya Akella

University of Wisconsin-Madison
{agember, ashok, akella}@cs.wisc.edu

Abstract. Handheld devices such as smartphones have become a major platform for accessing Internet services. The small, mobile nature of these devices results in a unique mix of network usage. Other studies have used Wi-Fi and 3G wireless traces to analyze session, mobility, and performance characteristics for handheld devices. We complement these studies through our unique study of the differences in the content and flow characteristics of handheld versus non-handheld traffic. We analyze packet traces from two separate campus wireless networks, with 3 days of traffic for 32,278 unique devices. Trends for handhelds include low UDP usage, high volumes of HTTP traffic, and a greater proportion of video traffic. Our observations can inform network management and mobile system design.

1 Introduction

Handheld devices—smartphones, portable music players, etc.—are quickly augmenting, and sometimes even replacing, laptops as the computing and Internet perusal platform of choice for users on the go. A 2009 EDUCAUSE study of technology on college campuses found 51% of undergraduates own an Internet-capable handheld and 12% plan to purchase one within the next 12 months [15]. A PEW study comparing 2007 and 2009 wireless Internet usage found a 73% increase in the rate Americans went online with their handhelds [10]. While the number of non-handheld portables, e.g. laptops, is also growing, usage of handheld devices is growing at a much faster pace.

In this paper, we seek to understand *how Wi-Fi traffic from handheld devices differs from non-handheld wireless clients, and what happens when handhelds override campus Wi-Fi networks*. Although many handheld users have cellular data plans, 802.11 Wi-Fi is still a preferred Internet access mechanism, when available, because of its higher bandwidth, lower latency, and lower energy usage. For our study, we use network traffic traces gathered from two independently-managed multi-AP campus wireless networks over a 3 day period. The traces have 32,278 unique clients, with 15% being handhelds.

We conduct an in-depth study of the *content and flow properties* of Wi-Fi handheld traffic. We examine transport and application protocols used, flow lengths and durations, and properties of content perused, e.g. prevalence of multimedia content and its nature and similarity in the content accessed by different users. We ignore low-level transmission, connectivity, and mobility issues as these have already been well studied [3, 8, 12]. To the best of our knowledge, these aspects of the differences between handhelds and non-handhelds have not been considered by prior studies. We believe that our examination of these issues is useful in informing future research on optimizing the performance of handheld devices operating in Wi-Fi networks. Specifically, our observations can help determine whether adopting prior approaches designed for non-handheld

devices, such as those for caching, content distribution and battery life savings, are applicable or not. Our study can also inform management practices for campus networks, e.g., network-wide Class-of-Service definitions for Multimedia traffic.

Compared to non-handheld wireless users, handheld users access a different mix of Internet services and content. Applications like web browsers and email clients are used on both types of devices, but content providers tailor content differently based on device type. Furthermore, the interface on handhelds in itself places limitations on the range of Internet-based and local network-based services users can access. Thus, the network traffic of handhelds is likely to differ in several crucial respects from non-handheld devices. The goal of this paper is to quantify the extent of these differences and identify the sources of the differences, where possible.

We present a broad collection of measurement insights. Our key findings are as follows. The majority of handheld traffic (97%) is web, with small amounts of email traffic. In contrast, 82% of non-handheld traffic is web, with miscellaneous UDP traffic (14%) accounting for most of the remaining share. Handhelds tend to have smaller TCP flows and a narrower range of flow durations. However, both types of devices have similar TCP flow rates, with a median rate of 0.8 Mbps. Looking in-depth at HTTP traffic, we observe that handhelds access content from a narrower range of hosts. However, we see equivalent amounts of similarity in content accessed by the same user for both device types. The top content type for handhelds is video, accounting for 40% of handheld traffic versus 17% for non-handhelds. Streaming video flows represent the largest, fastest, and highest throughput flows of all handheld flows.

2 Methodology

We collect and analyze data from two independently-managed campus wireless networks (Net1 and Net2). Full packet traces were captured from about 1,920 APs in Net1 over a period of 3 days during April 2010, yielding 8 TB worth of data. From Net2, full packet traces were captured from 23 APs for a period of 3 days in June 2010, yielding 50 GB worth of data. As an artifact of our collection method, we do not include traffic sent between wireless clients. However, we expect inter-client traffic is rare.

The packet traces contain data from all wireless clients connected to the network—laptops, smartphones, and other devices. Since we focus on the differences between handheld and non-handheld devices, we need to differentiate traffic based on device type. We rely on user-agent strings in HTTP packets as the primary method for differentiation. We identify handheld user-agents using a keyword list based on common knowledge and published lists [18].¹ Organizationally Unique Identifiers (OUIs) contained within device MAC addresses are used to confirm our device classifications. For the devices that do not send any HTTP packets, we use the OUIs of already classified devices to attempt classification based on OUI. Some devices (14%) remain uncategorized because their user-agent strings contain keywords associated with both types of

¹ **Handheld keywords:** Android, ARCHOS, BlackBerry, CUPCAKE, FacebookTouch, iPad, iPhone, iPod, Kindle, LG, Links, Linux armv6l, Linux armv7l, Maemo, Minimo, Mobile Safari, Nokia, Opera Mini, Opera Mobi, PalmSource, PlayStation, SAMSUNG, Symbian, SymbOS, webOS, Windows CE, Windows Mobile, Zaurus. See [7] for non-handheld keywords.

devices, or they send no HTTP traffic and their OUI is registered to a manufacturer that makes both types of devices; we exclude these devices from our analysis.

Over the 3 day capture periods, 32,166 unique clients connect to Net1 and 112 clients connect to Net2. Table 1 lists the number of clients of each type present in the trace data. Non-handheld devices account for the majority of clients in both networks. However, network admins provided anecdotal evidence that handhelds are much more prevalent than in the past, and industry and campus studies show the number of handhelds is expected to continue increasing [15]. We see handhelds from 7 primary vendors, with Apple iPods, iPhones, and iPads accounting for over two-thirds of all handhelds.

Device Type	Net1	Net2
Handheld	5060	9
Non-handheld	22485	90
Unknown	4621	13
Total	32166	112

Table 1: Client counts by device type

Protocol	Net1		Net2	
	Handheld	Non-hand	Handheld	Non-hand
UDP	5.9%	25.7%	4.5%	18.4%
TCP	92.0%	74.0%	93.0%	81.4%
IPsec	0.3%	0.05%	–	0.05%
Other	1.8%	0.35%	2.5%	0.15%

Table 2: Protocol usage (% of packets)

3 Protocols and Services

The protocols and services used by devices impact the performance of both the device and the enterprise wired and wireless networks. Different protocols and services respond differently to bandwidth limitations and congestion and contribute flows of varying sizes, durations, and frequencies to the overall traffic mix. Protocol mix also tells operators the mechanisms they must put in place to secure and monitor their networks.

3.1 Protocols

Network and Transport Protocols: At the highest level, we categorize traffic based on network and transport layer protocols (Table 2). As expected, the majority of traffic is TCP or UDP; the remaining traffic is IPsec (encrypted IP traffic) or network control traffic (ICMP, ARP, etc.). A major difference in protocol usage between the two types of devices is the amount of UDP traffic. Over 4x as many non-handheld packets are UDP compared to handhelds. In the presence of congestion, handhelds will use a fairer-share of bandwidth, versus non-handhelds which use more congestion-unaware UDP.

Application Protocols: We identify application protocols using Bro [13]. Table 3 shows the percentage of traffic in bytes for each category of application protocols. Web protocols account for the largest volume of traffic for both handheld (97% on Net1) and non-handheld (82% on Net1) devices. Almost one-third of Net2 handheld web traffic is HTTPS (versus 3% for Net1), but this is an artifact of a small sample size and a single large connection from one handheld. Email protocols are the second most popular application but account for less than 2% of traffic for both device types. We believe clients actually generate more email traffic than this, as shown by Falaki et. al for handhelds [6]; we attribute the low usage of email protocols to the common usage of web-based email and the potential for handhelds to simultaneously use 3G and Wi-Fi. Overall, our protocol usage observations are consistent with other studies [8].

Bro’s dynamic protocol detection can not identify the majority of UDP traffic for non-handhelds. More than 90% of the unidentified UDP traffic is large flows, from

Category	Protocols	Net1		Net2	
		Handheld	Non-hand	Handheld	Non-hand
Web	HTTP, HTTPS	97.0%	82.5%	91.1%	72.2%
Email	IMAP4, POP3, SMTP	1.51%	0.5%	—	0.04%
Chat	IRC	<	<	—	—
Remote	SSH, FTP	<	<	—	0.05%
Enterprise Services	IPP, LPD, NFS, SMB, LDAP, SQL	<	0.05%	—	0.3%
Management	DNS, NetBIOS, NTP, SNMP	0.2%	0.34%	1.52%	0.12%
Other TCP	<i>Unknown</i>	0.2%	2.9%	5.7%	8.7%
Other UDP	<i>Unknown</i>	1.0%	13.7%	1.7%	18.1%

Table 3: Application protocol usage by percent of bytes (< less than 0.01%, — none)

1 MB to 20 MB in size. While we don't know the exact nature of this traffic, we suspect that a majority of this traffic is likely from streaming media (e.g., Internet Radio).

As handheld usage in Wi-Fi networks continues to grow, HTTP traffic will become an increasingly dominant share of the traffic mix. Admins should consider deploying network middleboxes focused on HTTP traffic, e.g. in-network security scanners or web proxies, to better serve handheld security needs without impacting device efficiency.

3.2 TCP Flow Characteristics

We compare the TCP flow characteristics of handheld and non-handheld traffic to determine *if* and *how* flows differ between the device types. We look at the flow size, duration, and rate for the downlink half of TCP connections—data flowing from remote host to the wireless client—since the majority of data flows in this direction. Flows which do not end with a *FIN* or *RESET* are excluded. In all cases, the distributions for both Net1 and Net2 are equivalent; we omit the Net2 distributions for brevity.




Fig. 1: CDFs of TCP flow properties (Net1)

We observe that handhelds tend to have smaller flow sizes than non-handhelds. Figure 1a shows the median handheld flow size is 50 KB, versus 100 KB for non-handhelds. At the lower tail, there are fewer small flows for non-handhelds than handhelds; at the upper tail, the maximum flow size is larger for non-handhelds (2 GB) than handhelds (630 MB). The smaller handheld flow sizes are expected, as many content providers serve simpler or compressed content to mobile devices.

Handhelds and non-handhelds also differ in their distribution of flow duration. Figure 1b shows the median flow duration is approximately the same for both device types, but handhelds have a narrower range of flow durations. The middle 80% of handheld flows are 250ms to 15s long, compared to a range of 100ms to 75s for non-handhelds. The lack of long handheld flows can be attributed to typically short usage sessions [5]. We also look at flow durations for a few specific applications (full data in [7]). On average, web flows are 5x shorter for handhelds, which we suspect is caused by handhelds

being served simplified versions of many web pages. For email traffic, receiving protocols (IMAP, POP) have shorter average flows on non-handhelds, while the sending protocol (SMTP) has shorter flows on handhelds. We hypothesize the discrepancy in SMTP is caused by a higher likelihood of non-handheld users including attachments.

Downlink flow rates are shown in Figure 1c. Both device types have the same median rate of 10 Kbps, but only 10% of handheld flows are slower than 1 Kbps compared to 30% of non-handheld flows. Other factors associated with flow rate are consistent across both device types: (i) the average round trip time for 90% of TCP flows is less than 100 ms; (ii) only 4% of flows have one or more retransmissions due to retransmission time out, and 1% of flows have one or more retransmissions due to fast retransmit. Comparing duration and size of flows (not shown), we observe for both device types that small flows tend to have long durations, while large flows tend to have short durations.

4 Web Traffic

Web traffic accounts for almost all handheld data (97%) and a large fraction of non-handheld data (82%). HTTP is used so commonly because of its wide interoperability and support for many types of content. Web usage differs between device types because of differences in the way individuals use these devices. We see variation in (i) the range and type of hosts accessed and (ii) the type and length of content. We also observe that 82% of handheld HTTP traffic is consumed by non-browser applications, compared to 10% of non-handheld. Most notably, we see a higher usage of HTTP-based streaming media services on handhelds: video accounts for 42% of handheld HTTP content, versus only 23% for non-handhelds. Our analysis excludes partial HTTP streams (due to improper reassembly) and all data from the Net2 traces (due to anonymity concerns).

4.1 Hosts

HTTP hostnames, in combination with the type of content they provide, give a rough understanding of the types of services accessed by clients. Table 4 lists the top HTTP hosts based on the size (content-length) of the data served to the devices. We observe that handhelds access more multimedia content (by volume) than non-handhelds. Over 35% of handheld HTTP content originates from `googlevideo.com`, followed by 18% originating from `pandora.com`. Multimedia-type content is also the most frequent for eight of the top ten handheld hosts. In total, the top 10 handheld hosts account for 74% of handheld data, while the top 10 non-handheld hosts account for 42% of non-handheld data. These percentages indicate a much greater diversity in hosts for non-handheld devices. In addition, non-handheld devices are more likely to receive content from hosts providing more than text or multimedia content. For example, a Microsoft site hosting application downloads, `dlservice.microsoft.com`, appears in the top non-handheld hosts with `application/octet-stream` as the primary content type.

We also look at the top hosts based on number of HTTP requests (listed in [7]). The top 10 handheld hosts account for 30% of handheld HTTP requests, compared to 32% for non-handhelds. Also, there is a greater diversity of services in the top hosts by requests: social networking, streaming media, advertising, search, and news. In summary, both device types have a great diversity in the number of hosts they request data from, but handhelds receive most of their data (by volume) from a much smaller set of hosts.

4.2 Content Type and Length

The type of HTTP content accessed by devices further identifies the services used and highlights differences in traffic characteristics. We observe the largest volume of handheld content is video (42%), while images are the top type for non-handhelds (29%) (full data in [7]). Below, we discuss each of the MIME types in detail.

Bytes	Host	Top Content Types ³	Bytes	Host	Top Content Types ³
35.48%	googlevideo.com	v/mp4	11.45%	c.youtube.com	v/flv, v/mp4
18.12%	pandora.com	p/octet-stream, i/jpg	7.00%	pandora.com	p/octet-stream, i/jpg, a/mpeg
10.57%	phobos.apple.com	t/plain, i/jpg, v/mp4	6.63%	fbcfn.net	i/jpg, i/png, t/javascript
2.45%	fbcfn.net	i/jpg, t/javascript, i/png	4.63%	dlservice.microsoft.com	p/octet-stream
2.43%	vo.llnwd.net	v/mp4, a/mpeg	2.89%	vo.llnwd.net	v/wmv, a/mp4
1.23%	m.nbc.com	v/mp4, i/jpg, t/javascript	2.80%	stileproject.com	p/octet-stream, i/jpg, v/mp4
1.17%	espn.go.com	t/plain, t/html, i/jpg	2.53%	com.edgesuite.net	v/wmv, a/wma, p/octet-stream
1.16%	video.ted.com	v/mp4	1.69%	phobos.apple.com	t/plain, a/mp4, i/png
0.82%	gdata.youtube.com	t/atom+xml	1.51%	www.facebook.com	t/html, t/javascript
0.64%	s3.amazonaws.com	a/3gpp, i/jpg, i/png	0.94%	cdn.turner.com	t/javascript, i/jpg, v/flv

(a) Handheld

(b) Non-handheld

Table 4: Top HTTP hosts by response size (Net1)

Application content is data associated with specific applications, e.g. documents, compressed files, or streaming media. For both device types, *octet-stream*—a simple binary data stream—is the most common subtype, accounting for 86% of handheld and 51% of non-handheld *application* type data. The average *octet-stream* is 713 KB for handhelds ($\sigma = 882$ KB) and 189 KB for non-handhelds ($\sigma = 658$ KB). The second most common *application* subtype is RSS feeds for handhelds and Shockwave Flash for non-handhelds. No handhelds access Shockwave Flash content because these devices did not support Flash until very recently [1]. Over 185 different *application* subtypes are accessed by non-handhelds compared to only 58 subtypes for handhelds. This variety results from the greater diversity of applications running on non-handheld devices.

The content for regular web browsing falls mostly into the *image* and *text* content types. Three *image* subtypes—*gif*, *jpg*, and *png*—make up the majority of image content, with JPG images being the largest (average of 13 KB for handhelds and 11 KB for non-handhelds). HTML, CSS, JavaScript and XML are used for the web page itself. For both device types these *text* subtypes average 3-7 KB in length. Over two-thirds of the *text* content received by handhelds is *plain* text. This content is destined for non-browser applications retrieving data from a web service, e.g. a sports scores application.

The remaining MIME content types are multimedia traffic, namely *audio* and *video*. Multimedia accounts for 46% of handheld content and 29% of non-handheld content. In particular, *video* accounts for 93% of multimedia traffic in the handheld case and 80% in the non-handheld case. We examine video traffic in greater detail next.

4.3 Streaming Video

Streaming video is a major source of traffic for handheld devices. Video content accounts for 40% of all handheld traffic, compared to only 17% of all non-handheld traffic. We compare the flow characteristics of handheld streaming video, non-handheld video, and all handheld flows to understand the differences in handheld streaming media.

³ We abbreviate the MIME content types: v = video, a = audio, i = image, t = text, p = application.

As expected, handheld video flows are large compared to overall handheld traffic: 80% of video flows are > 50 KB in size, whereas 50 KB is the median among all handheld flows (Figure 2a). Nearly 20% of handheld video flows are > 1 MB in size, with a 400 KB median. Non-handheld video flows are even larger, with a 3 MB median.




Fig. 2: CDFs of video flow properties

Interestingly, handheld video flows appear to be of short duration. Figure 2b shows 80% of handheld video flows are less than 1 second in duration, with a median of 0.5 seconds. The median durations for all handheld flows and non-handheld video flows are significantly higher, at 5 and 50 seconds, respectively. Based on the short duration of handheld video flows, we expect high throughput rates. Figure 2c shows 80% of video flows have a rate faster than 0.8 Mbps, with a median of 2 Mbps. In contrast, the median flow rate for all handheld flows and non-handheld video flows is roughly 0.6 Mbps.

Overall, handheld video flows are long in size (although not as long as non-handheld video flows), significantly short in duration, and achieve high end-to-end throughputs which are comparable, if not slightly higher than non-handheld video flows. As handheld usage continues increasing, administrators should include Quality of Service mechanisms in their networks to support the video throughputs handhelds expect.

Video streamed to handheld devices differs from video streamed to non-handheld devices because of differences in decoding capabilities. Most streaming video services use Flash, but a lack of Flash support on handhelds results in MPEG 4 encoded content being served to them instead. In our traces, *mp4* (MPEG 4) is the top video type for handhelds and *flv* (Flash video) is the top type for non-handhelds. Video streaming sites like YouTube serve two versions: one encoded as *mp4* and the other encoded as *flv*.

We watch the same 3 minute video [2] from YouTube on both an Android HTC Dream smartphone and a laptop to measure the differences in video content served to the two different devices. On the phone we use the standalone YouTube application and on the laptop we use Mozilla Firefox. The handheld device receives 7362 KB `video/mp4`; the non-handheld device receives 11792 KB `video/flv`. Both versions have the same resolution of 320 x 240, but different encoding rates of 200 kbps and 231 kbps, for *mp4* and *flv* respectively. The audio is encoded at 128 Kbps for the *mp4* and 64 Kbps for the *flv*. The higher quality video is the *flv* and the higher quality audio is the *mp4*, but both versions are closely comparable. The main difference in the handheld content is a smaller size—about 62% of the size of the non-handheld version.

The median size of both handheld (316 KB) and non-handheld (1.7 MB) video flows are relatively small compared to the size of the sample video. In many cases, we observe videos being streamed over multiple sequential connections—due to connection resets—resulting in a few small flows for each video. However, by comparing the

combined size of these multiple flows to the size of the actual video, we observe that the size gap also results from handheld users watching only a fraction of most videos.

5 Content Similarity

In this section, we examine the similarity in the content perused by handhelds and compare it against non-handhelds. §4 focused on the *type* of content present in traffic; here we focus on the *bytes* that makeup the content. We evaluate the potential benefits of deploying a “chunk-based” content similarity suppression system, e.g. SET [14] or EndRE [4]. Eliminating duplicate chunks from network transfers by serving them from a local cache can improve the transfer throughput experienced by users and can help save mobile battery life by reducing network transmissions. Chunk-based schemes are more effective than object-based schemes, such as Web caches, as they are known to identify more duplicates, e.g., sub-object duplicates, uncacheable content, etc. Thus, our analysis places an upper bound on the benefits of using caching and similarity suppression.

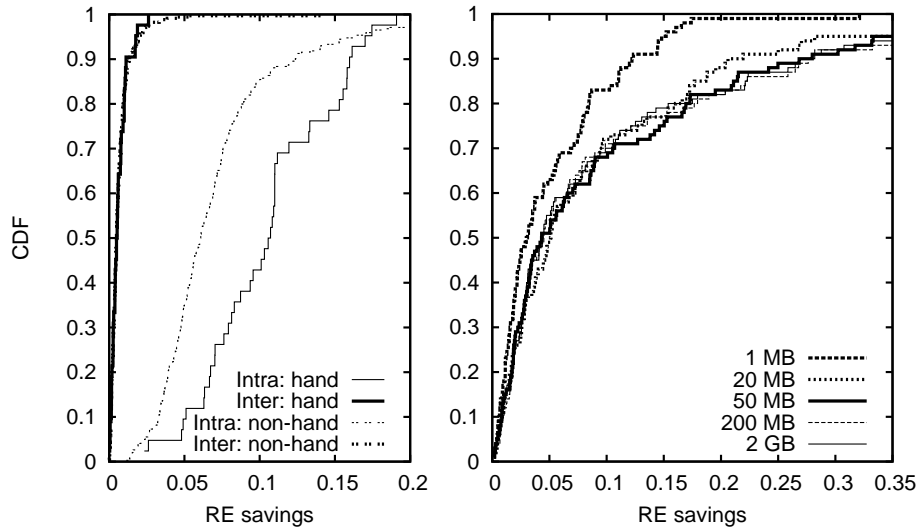
We identify two types of similarity: that found in content accessed by the same device (“intra-user”), and that found in content accessed by a different devices (“inter-user”). We divide packet payloads into chunks (32B to 64B in size) using value sampling [16]; then we determine if the chunks have appeared in an earlier access. Unless specified, we assume 2GB of chunks are stored across all users, as done previously [4].

In Figure 3a, we show the extent of intra- and inter-user content similarity observed over every 1 million packets (0.8-2GB) worth of handheld and non-handheld traffic. We measure average redundancy as the ratio of similar bytes to all bytes in 1 million-packet trace subsets. First, we observe a greater amount of similarity in handhelds than in non-handhelds. Second, similarity due to inter-user matches is quite small: less than 2% for > 95% of both handheld and non-handheld trace subsets. Third, we observe that in more than 40% of the non-handheld trace subsets, and more than 70% of the handheld trace subsets, $\geq 8\%$ of the similar bytes are due to intra-user matches. In some cases, we observed up to 20-25% intra-user similarity for both device types. Finally, the extent of intra-user similarity is greater for handhelds than non-handhelds.

Given that the dominant fraction of similar bytes belonged to intra-user traffic, we further delve into intra-user similarity. We explore the efficacy of deploying per device caches and the cache size configuration issues therein. We split the handheld traffic on a per device basis and study the effect of different dictionary sizes on amount of similarity identified per device. Figure 3b shows the CDF of similarity across the top 100 devices by traffic volume for different dictionary sizes. Almost 80% of users have less than 20% similarity with their own traffic. However, for certain users, the similarity proportion was much higher (more than 50%). Second, we observe that most of the similarities can be identified by using only 50 MB caches—larger caches exhibit diminishing returns. As handheld usage grows, admins should consider deploying per-device caching mechanisms to improve throughput and handheld energy savings.

6 Related Work

Multiple measurement studies have analyzed traffic patterns in campus wireless network. Hederson et. al identify session and application trends at Dartmouth College



(a) Average intrauser and interuser redundancy across multiple traces (b) Intrauser redundancy across top 100 (by bytes) handholds for varying dictionary sizes

Fig. 3: CDFs of redundancy

and observe how usage evolved four years later [8]. Wireless AP workloads at Dartmouth are compared to the University of North Carolina by Hernandez-Campos and Papadopouli [9]. Lastly, McNett and Voelker study the wireless access and mobility patterns of students using PDAs at UCSD [12]. While all of these studies focus on campus wireless networks, none explore in detail the applications used by handheld users and the traffic characteristics thereof. In addition, mobile device usage is a rapidly changing field and trends observed five years ago are different than today's usage.

More recent studies have focused on mobile device usage in public Wi-Fi, home Wi-Fi and 3G networks. Application, session, and mobility trends in the Google Wi-Fi network in Mountain View, CA were studied in 2008 [3]. The connections between geolocation and usage of specific types of web services was studied in an urban 3G network in 2009 [17]. In 2010, logs from 43 smartphones were analyzed to find commonly used application ports and properties of TCP transfers over a combination of 3G and Wi-Fi networks [6]. A second 2010 study analyzes the protocol usage and HTTP content size and types of handheld traffic extracted from DSL traces [11]. The 2010 studies are most similar to our work, but one focuses primarily on 3G traffic and neither looks in-depth at the multimedia content served to handhelds nor the redundancy in handheld traffic.

7 Conclusion

Handhelds have become a significant fraction of the client base in campus wireless networks, and their usage is expected to continue growing. Using traces from two separate multi-AP wireless networks, we identify key differences in the Wi-Fi content access and flow-level traffic characteristics of handheld and non-handheld devices. Our findings have potential implications for network management and mobile system design:

- 97% of handheld traffic is HTTP, allowing in-network security scanners to examine a single application protocol and provide significant security benefits for handhelds.

- Over twice as much handheld traffic is video, compared to non-handhelds, making Quality of Service mechanisms an important inclusion in network design.
- Lower HTTP host diversity and significant intra-user content similarity in handheld traffic, indicates per-device redundancy elimination systems can be beneficial.
- The smaller range of TCP handheld flow durations and the lower percentage of handheld flows with rates < 1 Kbps should be taken into account when designing wireless power save mechanisms for handhelds.

Network admins and mobile designers should take these observations into account when considering design and performance. The differences between handheld and non-handheld traffic will increasingly impact Wi-Fi networks as handheld usage grows.

References

1. Adobe announces availability of flash player 10.1 for mobile. <http://www.adobe.com/aboutadobe/pressroom/pressreleases/201006/06222010FlashPlayerAvailability.html>
2. Christmas lights gone wild. http://www.youtube.com/watch?v=rmgf60CI_ks
3. Afanasyev, M., Chen, T., Voelker, G.M., Snoeren, A.C.: Analysis of a mixed-use urban wifi network: when metropolitan becomes neapolitan. In: IMC. pp. 85–98 (2008)
4. Aggarwal, B., Akella, A., Anand, A., Chitnis, P., Muthukrishnan, C., Nair, A., Ramjee, R., Varghese, G.: EndRE: An End-System Redundancy Elimination Service for Enterprises. In: NSDI (2010)
5. Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R., Estrin, D.: Diversity in Smartphone Usage. In: MobiSys (2010)
6. Falaki, H., Lymberopoulos, D., Mahajan, R., Kandula, S., Estrin, D.: A first look at traffic on smartphones. In: IMC (2010)
7. Gember, A., Anand, A., Akella, A.: Handheld vs. non-handheld traffic: Implications for campus wifi networks. Tech. rep., University of Wisconsin-Madison (2010)
8. Henderson, T., Kotz, D., Abyzov, I.: The changing usage of a mature campus-wide wireless network. *Comput. Netw.* 52(14), 2690–2712 (2008)
9. Hernandez-Campos, F., Papadopouli, M.: A comparative measurement study the workload of wireless access points in campus networks. In: IEEE International Symposium on Personal, Indoor and Mobile Radio Communications. vol. 3, pp. 1776–1780 Vol. 3 (Sept 2005)
10. Horrigan, J.: Wireless internet use. Pew Internet & American Life Project (Jul 2009)
11. Maier, G., Schneider, F., Feldmann, A.: A first look at mobile hand-held device traffic. In: Passive and Active Measurement (2010)
12. McNett, M., Voelker, G.M.: Access and mobility of wireless pda users. *SIGMOBILE Mob. Comput. Commun. Rev.* 9(2), 40–55 (2005)
13. Paxson, V.: Bro: a system for detecting network intruders in real-time. In: USENIX Security Symposium. pp. 3–3. USENIX Association, Berkeley, CA, USA (1998)
14. Pucha, H., Andersen, D.G., Kaminsky, M.: Exploiting similarity for multi-source downloads using file handprints. In: NSDI (2007)
15. Smith, S., Salaway, G., Caruso, J.: The ECAR Study of Undergraduate Students and Information Technology. EDUCAUSE Center for Applied Research (2009)
16. Tolia, N., Kaminsky, M., Andersen, D.G., Patil, S.: An architecture for Internet data transfer. In: NSDI (2006)
17. Trestian, I., Ranjan, S., Kuzmanovic, A., Nucci, A.: Measuring serendipity: connecting people, locations and interests in a mobile 3g network. In: IMC. pp. 267–279 (2009)
18. ZyTrax: Mobile browser id strings. http://zytrax.com/tech/web/mobile_ids.html