

Characteristics of Mobile Web Content

Paul J. Timmins, Sean McCormick, Emmanuel Agu, Craig E. Wills

Department of Computer Science

Worcester Polytechnic Institute

Worcester, MA U.S.A.

Email: {ptimmins|mccorms}@wpi.edu, {emmanuel|cew}@cs.wpi.edu

Abstract—The World Wide Web is no longer tethered to our desktops and laptops. The Web has gone mobile, providing instant access to information anywhere and anytime. The mobile Web can be considered a shadow of the World Wide Web, implemented using specialized markup languages and design techniques adapted for comparatively limited mobile phones and PDAs. Despite the growing importance and usage of the mobile Web, surprising little is known about it.

This paper presents the results of a study of mobile Web content conducted in May and June of 2006. The study examines the content of over one-million mobile Web pages from around the world using a search-assisted crawling methodology to locate and study pages for three of the most popular mobile Web formats—WML 1.0, WML 2.0/XHTML Mobile Profile (XHTML-MP) and Compact HTML (C-HTML). The objective is to study the relative characteristics of these mobile Web content formats, as well as compare them with a similar sampling of non-mobile (HTML) content.

We found that WML is the dominant mobile Web content type, although regional differences do exist. We found that all three mobile content types studied were on the same order of magnitude for average page characteristics such as number of links (under 10) and number of images (around 1), but pages in the newest format, XHTML-MP, are 50% larger on average than those in WML. Not surprisingly, all of these characteristics are much smaller than for HTML content pages gathered with the same methodology. In terms of specific features, only 7% of pages used WML cards, but 50% of XHTML-MP servers dynamically adapted the content served based on the user agent. Finally, we found less than 4% of mobile pages contained ad objects, which is much less than for HTML pages.

I. INTRODUCTION

The World Wide Web is no longer tethered to our desktops and laptops. Web content is increasingly available in mobile Web formats that facilitate information access by cell phones, PDAs, Internet connected watches and other portable computing devices. Mobile users can now access sports, news, stock charts and other Web content while on the move.

Content targeted at mobile devices is typically designed to mitigate the lower bandwidths of wireless networks as well as the reduced CPU and storage limitations of mobile devices. In order to maintain reasonable download times, Web designers reduce the size of mobile Web pages and the number of images per page.

A number of early Web characterization studies informed and influenced the development of the wired Web [1], [2], [3]. As the mobile Web develops, it is important to understand the characteristics of how it is being used. This characterization shall inform optimal choices for configuring

network equipment and optimizing Web server parameters, such as buffer sizes and maximum number of incomplete TCP connections. PDA and cell phone manufacturers can also estimate minimum hardware specifications for future mobile devices. Understanding the nature of mobile Web content shall also drive more accurate simulations of Web content in the research community.

In spite of these benefits, surprisingly few attempts have been made to measure mobile Web content and quantify typical page and site characteristics. In this paper, we present the results of a large-scale study to characterize mobile Web content. The study examines the content of over one-million mobile Web pages collected during search-assisted crawls in May and June 2006. The study located and studied pages for three of the most popular mobile Web formats—WML, XHTML Mobile Profile (XHTML-MP) and C-HTML. We also use this same methodology to retrieve and measure wired Web content as a baseline of comparison.

We found that WML is the dominant mobile Web content type, although regional differences do exist—WML (WAP 1.0) is most popular in Europe and C-HTML (i-mode) is most popular in Japan. C-HTML crawling, and therefore C-HTML results, were significantly limited by the fact that many such sites are accessible only through NTT DoCoMo's "i-mode menu" service. We found that all three mobile content types studied were on the same order of magnitude for average page characteristics such as number of links (under 10) and number of images (around 1), but pages in the newest format, XHTML-MP, are 50% larger on average than those in WML. Not surprisingly, all of these characteristics are much smaller than for HTML content pages gathered with the same methodology. In terms of specific features, only 7% of pages used WML cards, but 50% of XHTML-MP servers dynamically adapted the content served based on the user agent. Finally, we found less than 4% of mobile pages contained ad objects, which is much less than for HTML pages.

This paper is organized as follows. Section II provides background on mobile Web technologies. Section III outlines the questions we intend to answer with this study and Section IV details the methodology used. The results of our study are presented in Section V with a summary of these results in Section VI. Section VII describes related work and Section VIII outlines potential future work as well as summarizes the findings of this work.

II. MOBILE WEB BACKGROUND

This section reviews technical details of three of the most popular mobile Web technologies (WAP 1.0, i-mode and WAP 2.0) and compares them to the wired Web and HTML. These technologies can be distinguished by the devices that support them, their protocol stacks, and their markup language.

WAP 1.0 was introduced in 1998 as the first Mobile Web standard [4]. Several companies including Nokia and Motorola teamed up to develop the initial Wireless Application Protocol (WAP) protocol stack. It was envisioned that WAP 1.0 would enable a wide range of devices including mobile phones, laptops and PDAs, to send email and access the Web. Due to the limited resources of mobile devices, a lightweight, XML-based standard called the Wireless Markup Language (WML), was developed. WML also supports a "deck of cards" feature that allows the Web programmer to aggregate multiple related pages (cards) into a batch (deck). WAP 1.0 is connection-oriented and a mobile user has to make a telephone call to the web server while web pages are being downloaded.

The i-mode (information mode) system was created in Japan at about the same time as WAP 1.0. It was deployed by NTT DoCoMo, the Japanese mobile network operator, in early 1999 and is loosely based on the WWW protocols. The i-mode system allows its users to email, surf the Web, and exchange images but requires specialized mobile handsets. Pages in i-mode are programmed using compact HTML (C-HTML) [5], a markup language that is similar to HTML 1.1. Most i-mode sites are accessible through NTT DoCoMo's "i-mode menu" service, which limits access to most i-mode sites to paying customers.

The second generation WAP 2.0 was introduced in 2001. It was designed by the WAP forum to be backwards compatible with the WAP 1.X protocols and WML. It includes a lot of the features of i-mode and Internet protocols, as well as new features. Compared with WAP 1.0, which has a maximum speed of 9.6 kbps, WAP 2.0 operates at speeds of up to 384 kbps. It supports XHTML, a new markup language that was developed for a variety of low computing power devices such as televisions, vending machines, mobile phones, PDAs, and watches. XHTML-MP [6], a mobile profile, extends XHTML basic by adding features to enhance the Web experience on resource constrained mobile devices. WAP 2.0 is designed to run over packet-switched networks and supports both push and pull models of content access.

III. STUDY

With this background, the broad goal of the study is to understand the characteristics of mobile Web content as it is currently being used. This broad goal encompasses a number of specific questions that form the basis for the methodology used in this work. These questions and their rationale are:

- 1) *Format usage*: What is the relative usage of the three markup languages—WML vs. C-HTML vs. XHTML-MP? The answer to this question establishes the relative use of these three formats by content providers.

- 2) *Geographic distribution*: What is the geographic distribution of content in the three markup languages? It is important to understand not only how much, but where the different content types are being used.
- 3) *Page sizes*: What is the distribution of markup (base page) and total page size for each of the three formats as well as baseline HTML content? A standard characterization for Web content is to understand the size of pages in terms of the number of objects they contain, the number of servers these objects come from and the total number of bytes contained.
- 4) *Page connectivity*: What is the degree of connectivity in terms of the number of links for mobile pages and are these links internal to the same domain or to different domains? This question examines how the level of connectedness compares amongst the three content formats and with HTML content.
- 5) *Image content*: What are the characteristics of image objects in mobile Web content in terms of number on a page and size? Images are commonly used in wired Web content. It is important to understand how much they are used in mobile content.
- 6) *WML cards*: Are unique schema features, such as WML cards, used? WML content can be served as bundles of pages or "decks." An interesting question is to understand how much this feature is used.
- 7) *User agent adaptation*: To what degree do servers adapt the markup type of the content based on the User-Agent field of the HTTP request header? This question affords better understanding on whether users need to explicitly identify the needed content type or whether servers can and do make the appropriate transformation.
- 8) *Advertisement content*: What is the presence of advertisement content in the mobile Web world? Previous work found that ad content exists on the majority of popular pages [7] and we are interested to understand its use in mobile pages.

IV. METHODOLOGY

This study was conducted in two phases. In the first phase, mobile Web servers were crawled to find and retrieve mobile Web content. An open source Web crawler was modified to classify content as mobile, non-mobile HTML, image or other. Mobile and non-mobile content was retrieved in its entirety, whereas only the size and URL of images were obtained.

The second phase of this study analyzed the retrieved pages to measure page size distributions, connectedness, and design features. Key features of each page were summarized in a MySQL database, allowing detailed analysis through SQL queries.

A. Mobile Content Crawler

To find and retrieve mobile Web pages, we modified Larbin, an open source Web crawler <http://larbin.sourceforge.net/index-eng.html>, to create a "Mobile Content Crawler." Larbin provides a configurable and

extensible framework for Web crawling, with many options to control the crawler’s behavior. The Mobile Content Crawler extends Larbin to identify mobile content, record page and image metadata to disk, (including HTTP response headers and content size), and retrieve the individual pages. The crawler was configured to use a 30-second delay between consecutive retrievals from a single server, with no delay for links to new servers. The effect of this is a preference on discovering new servers, but continued discovery of new pages within a site. Additionally, modifications were made to the Crawler so that it filtered non-mobile HTML content, recorded the image size/URL then discarded the image file, and ignored “robots.txt” (used by servers to prevent crawler access) so it could crawl search engine results.

In trial runs, Larbin was only modified to support mobile content, but not filter out non-mobile content, and seeded with a set of 14 starting mobile Web URLs, such as `mobile.espn.com` and `wap.yahoo.com`. These starting URLs were manually selected to represent a diverse population of content from a variety of mobile markup languages. During these trial runs, it was observed that a disproportionate number of HTML (non-mobile) Web pages were retrieved. As later results will show, this result is probably due to the higher connectivity (in terms of hyperlinks), of HTML content. The resulting effect was that retrieving HTML content reduced the number of mobile pages retrieved in that run.

To improve the crawler’s capability to retrieve the desired type of content, Larbin was modified to filter non-mobile content and retrieve only pages that could be explicitly identified as being mobile. This filtering was based on the content-type response header and the document type, if present. To reduce the volume of data stored, Larbin was also modified to first download images and store the size of the image in the header, using a preprocessor directive.

In addition to encountering problems in filtering out non-mobile content, the trial runs also showed problems in using a small fixed set of starting URLs. The result were not as diverse as desired either in the content type or the subject matter. As a consequence a search-assisted strategy was employed.

B. Search-Assisted Crawling

To address problems encountered in the trial runs, a search-assisted crawling strategy was employed for this work. It was noted that Google’s Mobile Web search engine (`mobile.google.com`) provides access to a large index of mobile web sites, and thus the search results as crawling starting points. This is similar to the strategy used in a previous study of Spyware [8]. Rather than directly select a set of starting URLs, the results of a Google Mobile Web search were used as crawling starting points, passing in specific keywords to the search engine. Google Mobile Web search allows searching of content by markup type (WML, XHTML-MP, or C-HTML). As a comparison, we also issued a search for HTML-based content using Google’s standard search engine.

A number of keywords were used to obtain search results for each of the four types of content to seed our crawler. These

keywords are shown in Table I, chosen to ensure diversity of search results. The upper portion of the table shows category-based keywords while the lower portion shows that we focused some searches on servers from specific country-based Top Level Domains (TLDs). These keywords are intended to obtain a broad set of pages for seeding.

TABLE I
SEARCH KEYWORDS USED FOR CRAWL SEEDING

news	sports	weather
games	portal	science
health	business	finance
arts	shopping	world
site:.jp	site:.uk	site:.ja
site:.au	site:.ve	site:.cn
site:.kp	site:.ca	

This study is based on four crawl runs done in May/June 2006, focusing on collecting HTML, WML, XHTML-MP/C-HTML, and C-HTML-only, all using the same keywords but different Google Mobile search restriction options. We found that, regardless of the search restriction option (ie: `mrestrict=wml`), crawl results are dominated by the most popular markup. Therefore, multiple runs were used, with later runs filtering out HTML and WML results. Additionally, the crawler was configured to report an appropriate User-Agent string in the HTTP request header, depending on the type of content we were attempting to crawl.

As with any crawling, the choice of starting points can bias the results. Search-engine assisted crawling, as used in [8], is biased by the search-engine’s results. Fewer than 15% of servers were directly linked from the search results, with the remaining servers being crawled indirectly by following hyperlinks. This high percentage of indirectly-crawled servers lessens the impact of bias caused by the search results.

An early goal of this research was to additionally contrast content based on subject matter, such as comparing news versus finance content. This goal evolved into the search-assisted crawling strategy, with the assumption that content could be characterize based on search keyword. However it was observed that a high degree of pages associated with multiple search keywords, thus this goal was set aside for future research. Alternative crawler strategies, including shallower searches, might yield results that are distinguishable by keywords.

C. Identifying Mobile Content

Two techniques were employed to identify content as being mobile content. First, the Document Type Declarations (DOCTYPEs) identify the DTD for a particular XML document. DOCTYPEs are optional, however were present in over 75% of servers crawled. Secondly, the HTTP CONTENT-TYPE response header identifies the expected type, such as HTML or image content. WML content is identified with a CONTENT-TYPE of “`text/vnd.wap.wml`”, and all other text content typically identified simply as “`text/html`”.

TABLE II
PAGE, SERVER AND DOMAIN CONTENT TYPE STATISTICS

Type	Num. Pages	Num. Servers	Num. Domains	Avg Pages/Server
WML	1,055,589	13,672	5,734	77
XHTML-MP	145,314	842	446	173
C-HTML	14,206	27	26	526
HTML	227,462	47,110	38,143	5

Content was first characterized by the CONTENT-TYPE, then by the DOCTYPE. The DOCTYPE of each page was used to classify content into the following categories,

- WML: DTD WML
- XHTML-MP: “XHTML Mobile Profile” or “XHTML Basic”
- C-HTML: “Compact HTML”

A possible limitation is our method of identifying content based on the DOCTYPE XML tag, which is only present in pages from 71% of servers. We observed the majority of the remaining pages contained “wap”, “imode”, “chtml”, “wml” or “mobile” in their URLs, but did not contain the necessary DOCTYPE. Manual inspections of these pages indicated that the presence of these keywords in a URL did not necessarily indicate mobile content, and therefore such content was excluded from the study.

D. Analyzing Page Content

Once the pages were retrieved and data about them stored in a MySQL database then the last phase of our study gathered data needed to answer the study questions. To obtain the size of a mobile Web page, queries were written to retrieve image and page sizes. This result indicates the amount of information that is downloaded by a mobile browser. From a resource perspective, large Web pages would consume excessive amounts of memory, CPU, battery power and wireless bandwidth. They would also take too long to download. However, small pages may not contain all the information that a user wants, making it necessary to establish new TCP connections to download additional pages.

We also examine the number of links on each Web page and distinguish whether the target pages are located on the same domain (internal link) or on a different domain (external link). Computing the internal versus external link numbers for both mobile and wired content is a means to compare their degrees of connectivity.

Links and images were counted uniquely per page, therefore multiple links to the same URL counted as one link. As well, the total page size was computed by adding the markup size to the size of each image in the page. Only complete pages, pages where the size of each image was measured, are reported so as to not skew the results, and images were counted only once per page as most browsers will not retrieve the same image more than once per page. In addition, total page size was normalized to compensate for the fact that pages with images were less likely to be completely retrieved by summing the weighted

average of the markup size of pages with no images with the markup and image size of pages with images.

V. RESULTS

This section presents results from our search-based crawling approach and subsequent analysis. The results are presented in a parallel order to the study questions posed in Section III.

A. Format Usage

Table II summarizes our crawl results concerning pages, servers and “domains” for each of the four formats. We define the domain of a server to be its 2nd-level domain¹ so servers such as `www.cnn.com` and `images.cnn.com` are each part of the `cnn.com` domain.

A key observation from Table II is that the number of WML pages found was an order of magnitude more than XHTML-MP, and two orders of magnitude greater than C-HTML. This result could either imply that there are indeed more WML pages in existence, or the results could also be skewed by the chosen search strategy. The low representation of C-HTML content is presumably due to NTT DoCoMo’s “i-mode menu” service, which provides paying customers access to pre-approved i-mode sites and thus is not accessible by the general public. Subsequent crawling runs using different search and filtering tactics were used to attempt to increase the number of XHTML-MP and C-HTML pages collected, including using Google’s international web servers and filtering out WML results. These subsequent crawling runs did not significantly increase the number of pages.

Table II also provides the average number of pages crawled per server. Recall from Section IV, the crawler was configured to wait at least 30 seconds before page retrievals from the same web site, and would immediately retrieve pages from newly identified servers. The resulting crawler behavior is to prefer breadth over depth. Thus, the HTML results are expected: an average of only 5 pages per server, showing that the crawler was continually identifying and crawling new servers and thus is expected from the highest degree of external links (Table VI). For mobile content, C-HTML and XHTML-MP had much higher numbers of pages per server than WML, a fact not explained by the number of external links per page. This results can be attributed to C-HTML and XHTML-MP pages linking to a less diverse set of servers than the WML pages.

¹In cases where the Top Level Domain (TLD) is a country code and the TLD is subdivided using recognizable domains such as “com” or “co” then the domain of a server is its 3rd-level domain.

B. Geographic Distribution

Tables III and IV summarize the percentage of unique Web servers by top level domain, which provides some means to understand the geographic distribution of the gathered pages. The international flavor of the results show that the usage of mobile content is global. As expected, WML is more popular in Europe and China where WAP is mostly used. The breakdown of C-HTML sites by top level domain is not reported, due to the significantly fewer sites included in the study.

TABLE III
TOP LEVEL DOMAIN BREAKDOWN FOR WML CONTENT

Domain	% of WML Servers
.com	30%
.ru (Russia)	22%
.cn (China)	13%
.net	8%
.hu (Hungary)	3%
.de (Germany)	2%
.org	2%
.cz (Czech Republic)	2%
.uk (United Kingdom)	2%
other	18%

TABLE IV
TOP LEVEL DOMAIN BREAKDOWN FOR XHTML-MP CONTENT

Domain	% of XHTML-MP Servers
.com	47%
.ru (Russian Federation)	15%
.net	10%
.jp (Japan)	4%
.de (Germany)	3%
.ch (China)	2%
.no (Norway)	2%
.cn (Canada)	2%
other	14%

C. Page Sizes

Figures 1 and 2 show the distribution of page sizes, in terms of markup alone and total page size, emphasizing the size difference between HTML and mobile content. Total page size was only reported for pages where all images were collected, so that partial page sizes were not reported.

Table V shows average sizes of mobile markup, as well as the average total (markup + images). Total size counts each unique image in a page exactly once, but background images are not included. WML markup objects are the smallest on average with 2,159 bytes. As a comparison, results published in 2003 report an average of size of 1,230 bytes [9]. XHTML-MP markup objects were larger than WML pages by 40% on average, at 3,018 bytes. C-HTML markup objects were close on average to those of XHTML-MP type at an average size of 2,911 bytes. As is expected, HTML markup objects are the largest, by an order of magnitude at an average size of 35,490 bytes. As shown in the last column of Table V, the relative results for the average total page size are comparable for all

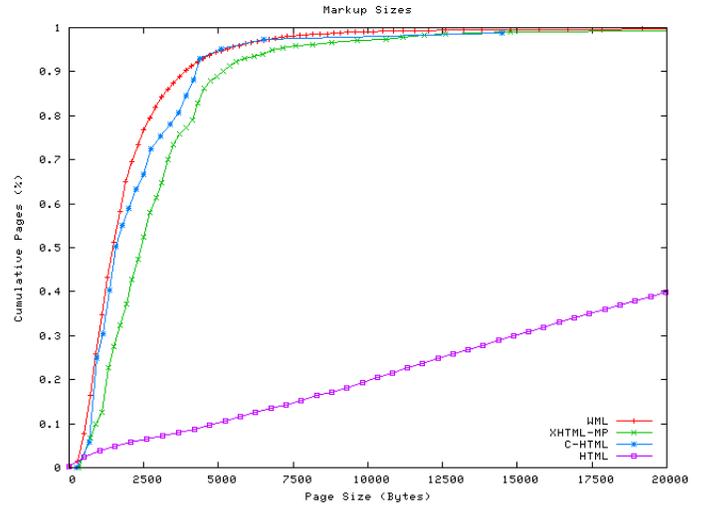


Fig. 1. Distribution of Content (Markup) Sizes

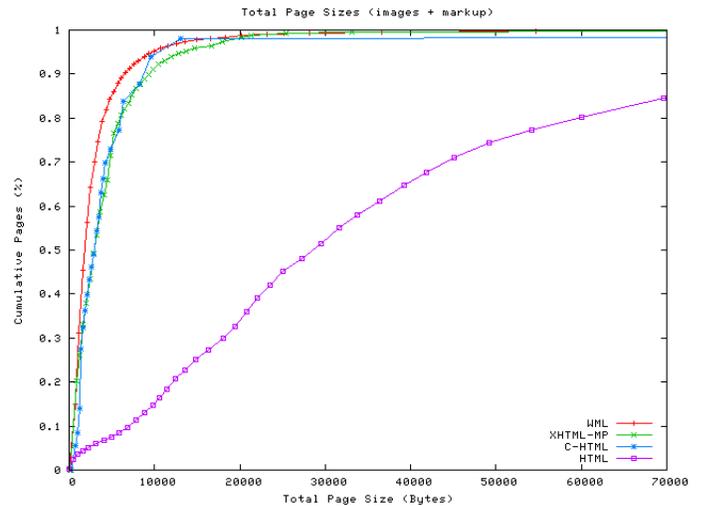


Fig. 2. Distribution of Total Page Size (Images + Markup)

four content types, although the sizes of XHTML-MP and C-HTML pages are more than 50% larger on average than WML pages.

Markup and image sizes are important considerations for content providers aiming to provide acceptable performance for their users. By providing content that is at or below average size, content providers can ensure that users will not suffer from abnormally long network transmission delays or memory consumption problems arising from large content. Mobile Web browsers and platforms should also be designed with expected content sizes in mind to ensure adequate memory and resources are provided to allow pages to be retrieved and cached. Results show that page sizes for the newer XHTML-MP format are an order of magnitude less than for HTML, but more than 50% larger than for WML.

TABLE V
CRAWLING RESULTS SUMMARY

Type	Avg Markup Size (Bytes)	Avg Total Size (Bytes)
WML	2,159	3,223
XHTML-MP	3,018	5,109
C-HTML	2,911	4,835
HTML	35,490	50,187

D. Page Connectivity

To examine the connectivity of pages, we measured the number of links from each page. Table VI shows the average number of internal and external links per page. Internal links are defined as links within the same domain, such as from `www.cnn.com` to `edition.cnn.com`. External links are defined as links across domains.

TABLE VI
AVERAGE NUMBER OF LINKS PER PAGE

Type	# Internal Links	# External Links	Link Density (bytes/link)
WML	7.6	1.1	248
XHTML-MP	8.3	1.1	321
C-HTML	6.2	2.6	331
HTML	56.7	24.2	439

Not surprisingly due to the significantly smaller page size, mobile content has significantly fewer links overall than HTML content. Figure 3 shows the median number of internal links for HTML pages is around 35 while for mobile pages the median is less than 5. Normalized by page sizes (see Table V), it is interesting to note that mobile markup had significantly higher link densities than non-mobile markup, with WML having nearly double the number of links per byte than HTML.

By necessity, mobile pages tend to be more concise than non-mobile due to the limited display sizes and user interface capabilities of mobile devices. Small displays require concise page design, to minimize excessive scrolling/paging of information. However, it is notable that at the high end, 10% of mobile pages had greater than 20 internal links, a surprisingly high number that would require scrolling by the user on most mobile devices. Figure 4 shows the three content types have a comparable distribution of external links, although a longer tail for C-HTML results in a higher average in Table VI.

E. Image Content

Table VII lists the average number of internal and external images per page while the distribution of the number of images per page is shown in Figure 5. A large number of pages contained no images (in `` tags). Our results found 44% of XHTML-MP pages, 61% of WML pages, and 70% of C-HTML pages contained no images. However, there is a long tail of pages that make use of more than five images. HTML pages, on the other hand, have significantly more images with a median of 12 images per page, and only 9% containing no images. Note that these results do not include background images, which are included through stylesheets. In addition, mobile content was less likely to contain external images,

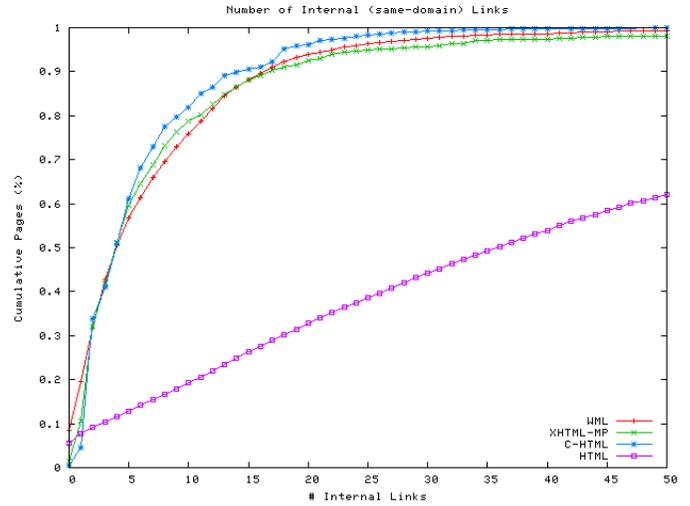


Fig. 3. Distribution of Internal Links Per Page

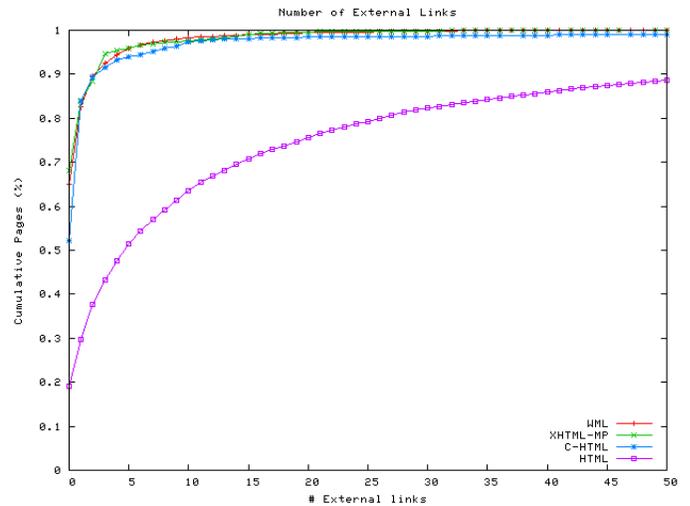


Fig. 4. Distribution of External Links Per Page

which was generally served from a fewer number of servers. The implication is that fewer TCP connections would need to be established, resulting in reduced bandwidth, power and overall page loading time.

Figure 6 shows the distribution of image sizes found for each content type. What is of interest in these results is that the median image sizes for each of the four content types is on the order of 1000 bytes. This result is because HTML pages use many images that are small in nature. As expected, HTML has more images that are larger.

F. WML Cards

The unit of transmission of WML content is a single deck, with each deck containing multiple individual pages which are referred to as cards. A card is the unit which is viewable in a Web browser. This feature allows a server to deliver multiple pages in a single transmission, eliminating a roundtrip to the

TABLE VII
AVERAGE NUMBER OF IMAGES PER PAGE

Type	# Internal Images	# External Images	Pages w/ No Images
WML	0.7	0.2	61%
XHTML-MP	1.3	0.3	44%
C-HTML	0.4	0.04	70%
HTML	11.9	3.7	9%

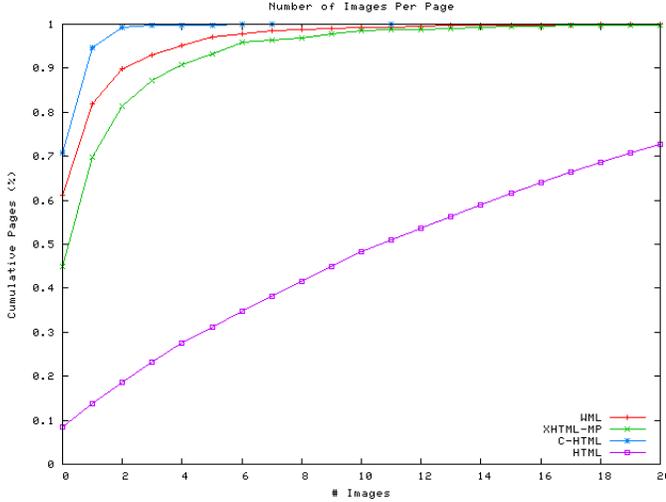


Fig. 5. Distribution of Number of Images Per Page

server in the event that a user accesses another page in the deck. This is a unique feature of WML 1.0.

As part of this study, we set out to understand whether WML cards were indeed used to improve user performance. Table VIII shows that 93% of WML pages crawled contained a single card, indicating rare usage of cards for improving user performance. Of the pages that contained more than a single card, 4% of pages contained exactly 2 cards. The remaining 3% of pages contained more than 2 cards. Of the pages containing more than one card: the mean number of cards was 16, and the standard deviation was 88.4.

From a server perspective, of the total of 13,672 unique WML servers, 15% contained at least one page with more than 1 card in the deck. There was a long tail of remaining servers, with 8% of servers containing more than 2 cards, 5% contained more than 3 cards, and 0.8% containing more than 10 cards.

There were a small, but nonetheless surprising, number of large outliers: over 2,000 pages from 16 servers had more than 100 cards in a deck and 83 pages from 2 servers contained more than 1,000 cards in a deck.

TABLE VIII
NUMBER OF CARDS PER WML PAGE

Avg # Cards Per Page	Max # Cards in a Page	Pages w/ 1 Card	Pages w/ 2 Cards
1.8	1,085	93%	4%

The lack of widespread use of cards to improve user

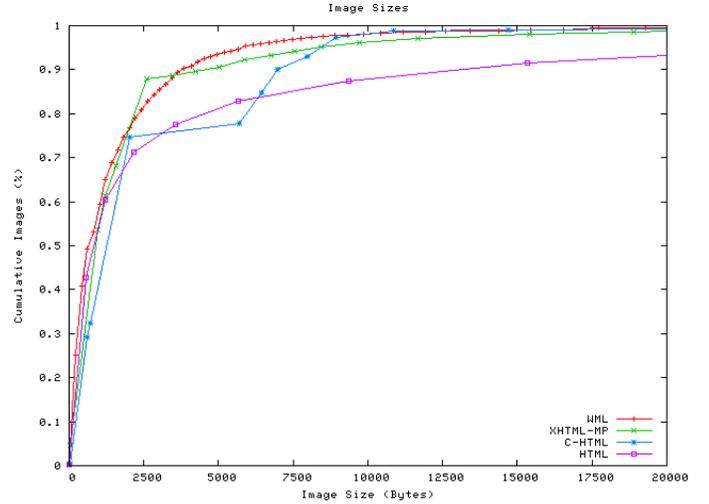


Fig. 6. Distribution of Image Sizes

perceived performance may be attributed to a number of causes. First, it is difficult to identify which content should be bundled with other pages. Second, the concept of cards and decks deviates from typical HTML Web design, requiring a shift by content providers. These results indicate that the lack of an equivalent feature (bundling pages into decks) in the XHTML-MP is not currently a significant drawback, due to the low usage of WML cards. Further, content providers are assured that despite the benefits of cards, few servers used this technique to improve the user experience.

G. User Agent Adaptation

Some servers were found to adapt the markup type of content depending on the client browser, detected through the User-Agent client header. To measure the extent to which this technique is used, we conducted tests in which we probed a given URL using six different User-Agent strings, shown in Table IX. This test was repeated for servers selected at random from our crawl results.

458 XHTML-MP servers were sampled, of which 50% (230) served different DOCTYPEs depending on the User-Agent. 98% (226) of these served WML content only to the WML User-Agent. In contrast, of the 5,291 WML servers sampled, only 0.5% (29) served different DOCTYPEs depending on the User-Agent.

This large difference may be based on a concern of backwards compatibility with older WML 1.0 browsers. As XHTML-MP is a newer generation markup language, not supported on all mobile browsers, the safest route is to serve WML content to older WML 1.0 devices.

H. Advertisement Content

Another characteristic we examined is the presence of advertisement content in mobile pages. Work in [7] found that ad content was contained on the majority of popular pages and we were interested in its use for mobile pages. To test for the presence of ads, we used the same methodology

TABLE IX
SELECTED USER AGENT STRINGS

Browser Type	User-Agent String
WML 1.0	EricssonR320/R1A UP.Link/4.1.0.1
XHTML-MP/WML 2.0	KDDI-KC31 UP.Browser/6.2.0.5 (GUI) MMP/2.0
i-mode	DoCoMo/2.0 SH901iCc100;TB;W24H12)
Internet Explorer 6.0	Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; SV1)
FireFox 1.5	Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.8.0.1) Gecko/20060111 Firefox/1.5.0.1
Palm Browser	Mozilla/4.76 (compatible; MSIE 6.0; U; Windows 95; PalmSource; PalmOS; WebPro; Tungsten Proxyless 1.1 320x320x16)

as [7] based on the Adblock Firefox extension [10], which blocks the retrieval of objects whose URL matches one or more pattern rules specified by the user. For rules, we use a ruleset named “Filterset.G” [11], which is commonly accepted as best practice for using Adblock to block ad content. For this analysis we used the 2006-03-08 Filterset.G ruleset version and converted the rules to Perl regular expressions (Perl and JavaScript use the same regular expression syntax) to identify all objects matching at least one rule. Results for the presence of ads for the four different content types we studied are shown in Table X.

TABLE X
PERCENTAGE OF PAGES WITH ADS FOR DIFFERENT CONTENT TYPES

Content Type	All	Top Level Domain						
		.com	.net	.org	.ru	.cn	.jp	.de
HTML	28.5	32.3	20.3	9.6	32.5	5.6	33.9	53.1
WML	3.6	4.8	3.6	1.3	2.9	6.5	–	0.2
XHTML-MP	3.9	1.7	2.0	0.0	–	–	0.1	27.0
C-HTML	1.0	1.9	0.0	0.2	–	–	0.1	58.8

Looking at the HTML content type results in Table X as a baseline for pages gathered with the methodology of this study, we see that less than 30% of all pages of this type were identified to contain at least one ad object. This percentage is about half of that obtained for all popular pages in [7]. We believe this reduction is because the pages used in this study were not explicitly identified as popular and because the data gathering we used in this study did not interpret JavaScript contained on pages, which was found to be a significant contributor to the presence of ads in [7].

The remaining content types in Table X show a much smaller percentage of pages with ad content for all pages and for selected TLDs. Examination of those URLs marked as ads confirms that similar to HTML content, there are many ad objects identified because they contain phrases such as “ads” or “sponsors”. There are also ads in mobile content identified because they are provided by specific servers such as `imrworldwide.com`, but not to the same degree as for HTML content.

Using the `.com` domain as a basis for comparison, the highest presence of ads is for the WML content type. The latter

columns in Table X show country TLDs that are prevalent in Tables III and IV. These results show the highest percentage of ads for WML content is found on Chinese (`.cn`) pages, while little presence of ads is found for the other two content types, save for a strong presence of ads in a relatively smaller number of German (`.de`) pages. Overall, we believe a number of factors contribute to the smaller percentage of ads found in mobile pages: 1) these pages contain fewer objects in general and hence are less likely to include ads as separate objects; 2) the ad filter rules we used were developed for non-mobile content and may not work as well for mobile content or for content from different countries; and 3) the presence of ads is not yet a priority for content providers and users, but just as with non-mobile content this will become an issue for both segments as the mobile Web continues to grow.

VI. SUMMARY OF RESULTS

The search-assisted crawling methodology employed in this work enabled us to obtain a number of notable mobile Web characterization results. These results and their significance are summarized in the following.

- *Format usage:* WML was the most encountered mobile content type by an order of magnitude. This result indicates that despite the emergence of new formats, there is still a large installed base of WML content.
- *Geographic distribution:* There was geographic diversity in where mobile Web content was found with a large amount of WML content in Russia and China, C-HTML content was most prevalent in Japan and X-HTML widely used in Russia. The C-HTML results are not surprising given that i-mode originated in Japan, but the Russian results were not expected, particularly since Russia was not even explicitly included as one of the per-country keywords.
- *Page sizes:* Focusing on the size of mobile Web content, the pages of XHTML-MP and C-HTML types are 50% larger on average than pages of WML type. These results indicate that newer format types are leading to larger mobile pages.
- *Page connectivity:* In terms of page links, the three mobile content types were comparable for both internal and external links. Not surprisingly, all mobile content types had an order of magnitude fewer links than HTML content. However all mobile content types had a higher link density than HTML when compared against page size.
- *Image content:* Image results showed that mobile content types had few images, but that XHTML-MP had the highest average number of images and C-HTML had the largest images on average. The number of images is an order of magnitude less than that for HTML content, although at least the median size of images is comparable for mobile and non-mobile content.
- *WML cards:* The use of the WML card feature is largely non-existent with 93% of pages utilizing only one card. This result indicates that a seemingly useful schema

features, such as cards, does not always translate into actual use.

- *User agent adaptation*: The results show that less than one percent of WML servers were found to employ user agent adaptation, while 59% of XHTML-MP servers did so. This result indicates a significant improvement in adaptation with a newer content type.
- *Advertisement content*: The presence of ads in mobile content is not yet significant at 4% of pages—certainly not to the same extent as for non-mobile HTML content at 30%. It would be interesting to track this figure as the mobile Web continues to evolve.

VII. RELATED WORK

Numerous Web characterization studies, such as [1], [2], [3], have been undertaken to answer questions about HTML content and the wired Web. A previous study [12] used the Larbin Web crawler to identify and explore characteristics of streaming media in the World Wide Web. A survey of Web metrics encompassing various measurements is presented in [13].

In terms of mobile content, an early study [14] studied the behavior of mobile users, but provides little comparable information on the characteristics of mobile content. In other work, [15] assessed the network traffic patterns resulting from WAP traffic, but it focuses on a simulation of WAP traffic with no direct measurements to characterize the WML content.

An evaluation of the usability of various mobile Web page design decisions was performed in [16]. Another piece of work on usability was [9], which developed a Wireless Usability Software Agent (WUSA) to perform a retrieval-based evaluation of WAP servers. This study was more limited in the type of content it considered (only WML), the number of pages it analyzed (~100K) and the depth of its analysis for the pages retrieved.

VIII. CONCLUSIONS

In this work we have presented results from an extensive characterization study of mobile Web content in a variety of content formats from Web servers around the world. An important piece of future work in such a characterization study is to repeat measurements at periodic times with the same methodology to understand trends that are occurring with the use of mobile Web content.

There is also other work to be done. One interesting question to answer is whether the mobile Web presentation depends on the device that is used. Answering this question requires separate, but near-time crawls of Web servers and making comparisons of the resulting information. Another direction of future work is to extend the set of metrics used to include others defined by [13]. These metrics include servers compactness, using degree distributions to see if Lotka's Law holds for the mobile Web, and the Small-World Model to see if the average distance between any pair of nodes is a small number. The presence of usability characteristics [17], such as horizontal scrolling versus vertical scrolling versus pagination

as well as the use of multiple pages versus a single large page could also be studied.

In the course of this project, it was discovered that content designed for Palm OS devices were observed to contain a specific META tag, but the DOCTYPE and CONTENT-TYPE fields of this content indicated it as HTML. We did not research this issue further in this work, but plan to do so in future work.

In summarizing this work, the search-assisted crawling methodology we employed enabled us to locate and study over one-million Web pages from around the world for WML, XHTML-MP and C-HTML mobile Web content. We found that WML is the dominant content type, although regional differences do exist.

We found that all three content types were on the same order of magnitude for page characteristics such as number of links and number of images, but pages of the newest format, XHTML-MP, are 50% larger than WML. Not surprisingly, all of these characteristics are much smaller than for HTML content pages gathered with the same methodology.

Looking for specific features, we found that few WML pages utilized the "deck of card" feature. We also found that almost no WML servers employed user agent adaptation. While most XHTML-MP servers did so. Finally, ad objects currently do not have a big presence in mobile Web pages.

REFERENCES

- [1] M. Arlitt and C. Williamson, "Internet web servers: Workload characterizations and implications," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 631–645, Oct. 1997.
- [2] J. E. Pitkow, "Summary of WWW characterizations," in *Proceedings of the Seventh International World Wide Web Conference*, Brisbane, Australia, Apr. 1998.
- [3] P. Barford and M. Crovella, "Generating representative web workloads for network and server performance evaluation," in *Proceedings of the ACM SIGMETRICS '98 Conference*. ACM, June 1998.
- [4] W. Forum, "Wireless application protocol: Wireless markup language specification," 1998. [Online]. Available: <http://www1.wapforum.org/tech/terms.asp?doc=wml-30-apr-98.pdf>
- [5] W. Consortium, "Compact html for small information appliances," 1998. [Online]. Available: <http://www.w3.org/TR/1998/NOTE-compactHTML-19980209/>
- [6] O. M. Alliance, "Xhtml mobile profile," 2001. [Online]. Available: <http://www.openmobilealliance.org/tech/affiliates/wap/wap-277-xhtmlmp-20011029-a.pdf>
- [7] B. Krishnamurthy and C. Wills, "Cat and mouse: Content delivery tradeoffs in web access," in *Proceedings of the International World Wide Web Conference*, Edinburgh, Scotland, May 2006.
- [8] A. Moshchuk, T. Bragin, S. D. Gribble, and H. M. Levy, "A crawler-based study of spyware in the web." in *NDSS*. The Internet Society, 2006.
- [9] T. Ceska, S. Kuhlins, and H. Nosekabel, "Evaluation of wireless usability with a java software agent," in *Proceedings of the 7th IASTED International Conference on Internet and Multimedia Systems and Applications*, Honolulu, Hawaii, August 2003.
- [10] "Adblock," <http://adblock.mozdev.org/>.
- [11] "Official home of Filterset.G," <http://www.pierceive.com/>.
- [12] M. Li, M. Claypool, R. Kinicki, and J. Nichols, "Characteristics of streaming media stored on the web," *ACM Trans. Inter. Tech.*, vol. 5, no. 4, pp. 601–626, 2005.
- [13] D. Dhyani, W. K. Ng, and S. S. Bhowmick, "A survey of web metrics," *ACM Comput. Surv.*, vol. 34, no. 4, pp. 469–503, 2002.
- [14] A. Adya, P. Bahl, and L. Qiu, "Analyzing the browse patterns of mobile clients," in *Proceedings of the ACM SIGCOMM Internet Measurement Workshop*, San Francisco, CA, Nov. 2001.

- [15] I. C. Y. Ma and J. Irvine, "Characteristics of wap traffic," *Wirel. Netw.*, vol. 10, no. 1, pp. 71–81, 2004.
- [16] L. Chittaro and P. D. Cin, "Evaluating interface design choices on wap phones: Navigation and selection," *Personal Ubiquitous Comput.*, vol. 6, no. 4, pp. 237–244, 2002.
- [17] G. Buchanan, S. Farrant, M. Jones, H. Thimbleby, G. Marsden, and M. Pazzani, "Improving mobile internet usability," in *WWW '01: Proceedings of the 10th international conference on World Wide Web*. New York, NY, USA: ACM Press, 2001, pp. 673–680.