A View from the Other Side: Understanding Mobile Phone Characteristics in the Developing World

Sohaib Ahmad
LUMS
Abdul Lateef Haamid
LUMS
Zafar Ayyub Qazi
UC Berkeley
Zhenyu Zhou
Duke University
Theophilus Benson
Duke University
Ihsan Ayyub Qazi
LUMS

ABSTRACT

Mobile devices are becoming increasingly dominant in the developing world. However, there is little insight into the characteristics of devices being used in such regions. Using a dataset of ~0.5 million subscribers from one of the largest cellular operators in Pakistan, we analyze the characteristics of cell phones based on different features (e.g., CPU, memory, and cellular interface). We identify potential device-level bottlenecks for Internet access and analyze the security implications of the phones being used. To aid the analysis of cell phones, we propose abstractions (e.g., connectivity, capacity, and device security) and cluster phones based on these abstractions. Our analysis reveals interesting insights for improving mobile web performance.

Keywords

Cellular Networks; Mobile Devices; Developing Regions

1 Introduction

Mobile devices are becoming increasingly common in developing regions. The International Telecommunication Union (ITU) reported that mobile-broadband subscriptions reached 2.3 billion by the end of 2014, with 55% of them in developing countries [19]. Mobile devices are being used for a variety of services in such countries including agricultural information dissemination, education, and health care delivery [23, 33].

While the Internet infrastructure has been steadily improving in developing regions and several efforts (e.g., Google's Project Loon [15] and Facebook's Connectivity Lab [3]) focus on providing Internet connectivity via satellites, balloons, drones, and planes, anecdotal evidence suggests that the common use of low end devices with a slow Internet connection can lead to poor user experience in developing regions [6]. Recent studies have explored the issue of poor application performance from a variety of angles, ranging from core infrastructural issues to geographical locations [31, 37]. For example, recent studies found that the lack of good caching infrastructure and DNS servers are the primary causes of poor performance [21, 37]. Another related study has shown CDN server placements and routing protocols as primary performance issues [32].

However, there is very little information available about the characteristics of devices being used in developing countries and how they may impact Internet access. This paper presents an in-depth study of cell phones from one of the largest cellular providers in Pakistan. Using a dataset of ~0.5 million subscribers, we, (a) present a taxonomy of cell phones based on different features (e.g., CPU, memory, OS type, and WiFi support), (b) discuss the implications of cell phone characteristics (e.g., maximum data rates, user experience due to browser support, and security vulnerabilities) for Internet access and identify potential performance bottlenecks, and (c) propose abstractions for classifying phones along different axes (e.g., connectivity, capacity, and flexibility).

Unlike prior works (e.g., [27]) that study the adoption and usage of smart phones, our unique focus on the individual features of cell phones reveals new insights, showing how mobile devices might be a major bottleneck to Internet access. This in turn can inform various stakeholders (e.g., content providers and Internet service providers) in improving mobile web performance. We believe this study provides a missing piece of the puzzle to understand Internet access in developing regions.

We make the following observations:

- We find that 66% of the cell phones support only GSM (52%) and GPRS (14%), thus they can expect to achieve no more than 40 kbps of data rates[1]. On the other hand, only 11.2% of the cell phones had HSDPA (a.k.a. 3G/3G+) or LTE capable handsets.
- We find that ~30% cell phones have CPU speeds of less than 500 MHz. The distribution of memory sizes was found

[1]This is a fairly common GPRS data rate in Pakistan.

IMC 2016, November 14-16, 2016, Santa Monica, CA, USA

ISBN 978-1-4503-4526-2/16/11...$15.00
DOI: http://dx.doi.org/10.1145/2987443.2987470

to be heavy tailed with most phones (57%) having less than 10 MB of memory and a small fraction (0.48%) having more than 1 GB of RAM. Such phones can often lead to poor performance for web browsing and web-based multimedia applications [31, 34].

- Our study shows that 58.1% of the phones being used have actually been discontinued by the vendors. As such, these phones may be vulnerable to well-known security attacks. Although, most phones use vendor-specific operating systems, Android and Symbian were the most popular mobile operating systems. Android 2.3, 4.0, and 4.1 were the most widely used versions, all with well-known security vulnerabilities [4, 14].

- We find that a significant fraction of phones (49.5%) in our dataset support only wireless application protocol (WAP) browsers. Unlike modern web browsers (e.g., Safari or Chrome), a WAP browser only supports a stripped down version of XHTML and uses the WAP protocol for accessing the Internet. However, it does not support JavaScript.

These findings suggest that device-level bottlenecks may be common in developing regions and should be taken into consideration when designing services for users in such regions. These insights can inform content providers, service providers, and even vendors on improving access to services. For example, we find that although CPU speeds are high in several phones, they are equipped with very small memory sizes. Such phones can benefit by having specialised web proxies that preprocess websites performing memory intensive computations.

Our work highlights the fact that simply changing the cellular network technology is insufficient, changes in the network must be aware of device limitations. To this end, our work reinforces the importance of developing regions research that deals with building low cost data communication channels (e.g., using SMS or voice as the transport mechanism) [25], building specialized proxies for developing countries [35], and designing applications suited for low-end feature phones [16, 28, 30]. In particular, using data channels over SMS/voice can enable or augment existing low data rate services (e.g., GSM and GPRS) to further improve the actual data rates. Moreover, non-data channels (e.g., voice) can also be directly used for building social networking, education, and health care applications. For example, Polly [30], a viral telephone-based system aims to reach low-literate population for development related services through a voice-based game.

In summary, this paper makes the following key contributions:

- We present a data-driven study of the characteristics of mobile devices in a developing country. Towards this end, we build a custom cell phone database containing 19 cell phone features (§2).

- We analyze the security implications of the cell phones being used, identify potential device-level bottlenecks and discuss how they may impact Internet access in a developing area like Pakistan (§3).

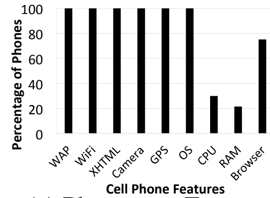

(a) Phones vs. Features

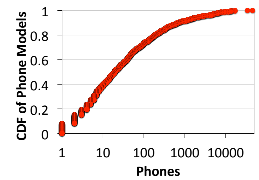

(b) CDF

Figure 1: (a) Fraction of cell phones whose feature information was available in our database and (b) CDF of the number of phones for different phone models.

- We develop a tool to characterize website support for WAP browsers. Using this tool, we survey the top 300 websites in Pakistan and found that while 25% of websites provide a significantly smaller version of their site for WAP browsers, only 6% are truly WAP compliant (i.e., no JavaScript or rich content) (§4).

- To aid the analysis of cell phones, we propose abstractions (e.g., connectivity, capacity, flexibility, and device security) and cluster phones based on these abstractions. Our analysis reveals interesting insights for improving mobile web performance (§5).

2 Dataset

Our dataset comprises twelve months (Jan 2014 – Dec 2014) of anonymized Call Detail Records (CDR) for district Jhelum from one of the largest cellular operators in Pakistan. The Jhelum district has an area of approximately 3,500 km^2, and a 2013 population estimate of ∼1.2 million. The dataset contains information about ∼0.5 million unique subscribers and the mobile devices they used. As 3G and 4G services were launched in Pakistan in April, 2014, we present analysis only for the month of December, 2014 to capture the latest set of cell phones being used.

We use the available device information (i.e., cell phone name and model) to construct a cell phone database by fetching information from online sources [8, 10]. Our dataset contains 19 cell phone features including information about the cellular interface supported (e.g., GPRS, EDGE, HSPDA, and LTE), WiFi interface (e.g., 802.11b, 802.11g, and 802.11n), CPU speed, amount of memory, operating system (OS), phone status (i.e., continued/discontinued), camera, and GPS, email, and browser support. Figure 1a shows the fraction of cell phones as a function of different features whose information was available to us[2]. For example, our database had information about the browser used in 77.3% of the cell phones.

Our dataset contains ∼4000 unique cell phone models. Figure 1b shows the cumulative distribution function (CDF) of the number of phones having a phone model. Observe that ∼80% of the phone models have less than 200 phones

[2]Due to space limitations, Figure 1a shows information for only a subset of the features.

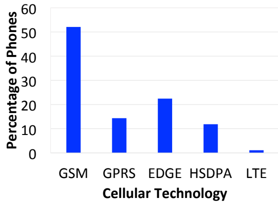

(a) Phones

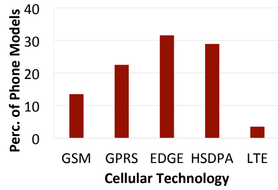

(b) Phone Models

Figure 2: Distribution of cellular technology across phones and phone models.

whereas only ∼5% of the phone models have at least 5000 phones each.

3 Taxonomy

In this section, we present a taxonomy of cell phones in our dataset based on a variety of features including cellular/WiFi interfaces, CPU speed and RAM capacity, OS type and phone status, and browser support.

3.1 Cellular/WiFi Interfaces

Cellular Interface: We now present a classification of phones based on the cellular technology they support. We find that 52% of the phones support only GSM as shown in Figure 2. GPRS and EDGE are fairly common and together make up 36.8% of the phones. HSPDA and LTE make up 11.2% of all the phones in our dataset.

Table 1 shows the typical downlink data rates supported by these technologies. Given that most phones only support GSM, their achievable data rate is upper bounded by 9.6 kbps [12], which is too low for supporting most popular data services and applications. ∼66% of the phones (i.e., GSM and GPRS combined) can expect to achieve no more than 40 kbps of data rates (which is a typical data rate offered by GPRS [13]). Together with EDGE, 88.4% of the phones can expect no more than 236 kbps of data rates[3].

Cellular Technology	Data Rates (kbits/s)
GSM [12]	9.6
GPRS [7, 13]	40-171
EDGE [5]	120-384
HSDPA [9]	600-13400
LTE [11, 18]	3000-36000 (max: 300 Mbps)

Table 1: Typical downlink data rates for difference cellular technologies.

Cellular technologies across cell phones models: We also study the distribution of cellular technologies over the different cell phone models (∼4000) present in our dataset. We found that only 13.6% of the phone models in our dataset had GSM support as shown in Figure 2b but they made up 52% of the phones in our dataset. ∼29% of the phone models supported HSDPA (a.k.a. 3G+) but these represented only 10.6% of the phones in our dataset. A possible reason for this trend is the high cost of such cell phones.

[3] This is the typical EDGE data rate offered in Pakistan [17].

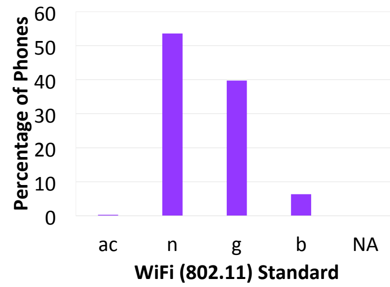

Figure 3: Distribution of WiFi standards across cell phones in our dataset.

WiFi Interface: In our dataset, only 12.5% of the phones had WiFi interfaces whereas 30.9% of the phones models had WiFi support. Figure 3 shows how different WiFi standards are distributed across phones that were equipped with WiFi interfaces. Interestingly, 53.6% of the phones had 802.11n capability and 39.7% had phones that only supported 802.11g. This in contrast to countries like the US, where 97.7% of devices were reported to be 802.11n capable during the same time period [22].

It is interesting to observe that despite the popularity of 802.11n, which supports a maximum data rate of 600 Mbps, a significant fraction of cell phone models (∼46.1%) are still limited to the low data rate 802.11b and 802.11g standards, which provide a data rate of up to 11 Mbps and 54 Mbps, respectively.

3.2 CPU/RAM

Cell phones with slow CPUs and small memory sizes often lead to poor application performance (e.g., for web browsing and web-based multimedia [28, 34]). Therefore, we now present a classification of phones based on their CPU speeds and memory sizes.

Processor: Table 2 shows the CPU speeds of cell phones in our database. We found that ∼32% of the phones had CPU speeds of less than 500 MHz. While cell phones with CPU speeds between 500 MHz and 1 GHz were most popular, cell phones with greater than 1 GHz speeds and multiple cores were much less popular. We found that 87% of the phone models with slow CPUs (i.e., less than 500 MHz) supported Email and 48% of such phones supported running third party applications, suggesting slow performance for CPU intensive tasks.

CPU Range (MHz)	Phones	% of Phones
0-500	19899	32%
500-1000	35647	57.4%
1000-1500	6333	10.2%
1500-2000	246	0.4%

Table 2: CPU speeds of cell phones in our dataset. The minimum and maximum CPU speeds were 104 MHz and 1900 MHz (8 cores), respectively.

Memory/RAM: Figure 3 shows the distribution of memory sizes of cell phones. We find that 57% of the cell phones had less than 10 MB of RAM whereas only 0.485% of the phones had more than 1 GB of memory. The minimum and

Mobile Operating System	Number of Known Security Vulnerabilities [14, 4]
Android Versions (1.5, 1.6, 2.0-2.3, 4.0-4.4, 5.1)	≥ 141
Windows Mobile/Phone (2003 SE, 6.0, 6.1, 6.5, 7, 7.5, 8)	≥ 9
Blackberry OS	≥ 8
iOS	≥ 428
Symbian (7.0, 8.1, 9.2-9.4)	≥ 3

Table 4: Some known security vulnerabilities (including DoS attacks, memory corruption, and code execution) in mobile operating systems found in our dataset.

Memory/RAM (MB)	Phones	% of Phones
0-5	31506	32.1%
5-10	24381	24.9%
10-128	10762	11%
128-256	4974	5.1%
256-512	17167	17.5%
512-1024	8824	9%
1024-2048	448	0.46%
2048-4096	25	0.025%

Table 3: Memory sizes for cell phones in our dataset.

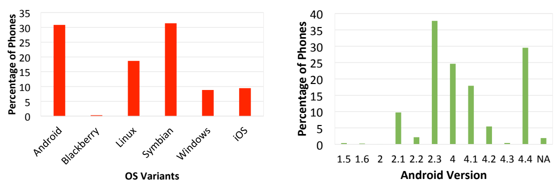

(a) OS Variants (b) Android Versions

Figure 4: Distribution of (a) operating systems across phones and (b) android versions across phones.

maximum memory sizes we observed were 1 MB and 3 GB, respectively. We found that 95% of the phones with small memory (i.e., $\leq$ 5 MB) supported Email and 84% of such phones supported running third party applications.

3.3 OS Type/Phone Status

We now discuss the distribution of mobile operating systems across phones. We found that 84.8% of the phones had vendor-specific operating systems whose information was not available. Figure 4a shows the distribution of cell phone OSes across the remaining 15.2% phones ($\sim$76k phones). Observe that 30.8% of the cell phones used Android followed by Symbian (31.4%).

We found that the three most popular versions of Android were 2.3, 4.0, and 4.1 as shown in Figure 4b. The observed mobile operating systems have well-known security vulnerabilities as shown in Table 4 [4, 14]. As such, these devices can be readily exploited to launch various attacks such as DoS attacks, attacks on privacy, and memory corruption.

Phone Status: Surprisingly, we find that 58.1% of the phones have been discontinued or cancelled by the vendors. These phones may be vulnerable to known security attacks as there is no continuous vendor support for addressing them. Such phones are likely to remain available in developing countries for quite sometime as they are low cost and reselling of phones is fairly common in such regions.

3.4 WAP Support

A WAP browser is designed for mobile devices and uses the WAP for Internet access. WAP 2.0 uses a cut down version of XHTML with end-to-end HTTP. Mobile devices process XHTML Mobile Profile (XHTMLMP), the markup language defined in WAP 2.0.

We found that 49.5% of the phones supported WAP whereas 89.8% of the phone models had WAP. Figure 5 shows the distribution of WAP versions found in our dataset. Observe that most phones supported WAP 2.0 (76.4%) whereas only 5.5% supported WAP 1.1. We expand on this in the next section and explore support for WAP from popular websites in Pakistan.

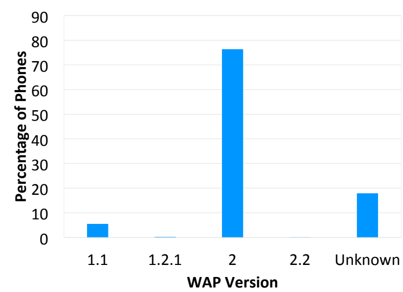

Figure 5: Distribution of phones which support WAP.

4 Characterizing Support for WAP Browser

In the last section, we observed that a significant fraction of phones (49.5%) in our dataset supported only WAP browsers and not modern browsers (e.g., Safari or Chrome). Recall that unlike modern rich browsers, WAP browsers support a stripped down version of XHTML. More importantly, WAP browsers do not support JavaScript and do not feature rich multimedia content (e.g., Flash and videos). Naturally, an interesting question is this:

"*Given that WAP is a dominant browser, what fraction of websites support WAP browsers?*"

To answer this question, we develop a web crawler that visits the top 300 websites in Pakistan according to the Alexa list [2]. The web crawler sets its "user-agent" to match those of the most popular devices in our dataset: Nokia3360.

We compared the website delivered when the crawler used a WAP user-agent against the website delivered when a Chrome user-agent was used. We observed that 25% of the sites explored delivered a lighter version (with fewer images and scripts) of their website to the web-crawler when it used a WAP specific user-agent. Even more surprisingly, we found that only 6% of the sites delivered truly WAP friendly versions with no JavaScript or rich multimedia content. Unsurprisingly, the 6% included large online service providers (e.g., Facebook, Google, Live, and Amazon) and several local sites, e.g., ummat.net (one of Pakistan's popular Urdu newspapers).

5 Characterizing Cell Phones

To aid analysis, we now propose four abstractions for cell phone features that represent groups of similar features[4]: connectivity, capacity, flexibility, and insecurity. These abstractions allow us to analyze the (a) network bandwidth offered by cell phones, (b) their computational resources, (c) flexibility in supporting third party applications and running modern browsers, and (d) vulnerability of a device to known attacks[5]. Moreover, by reducing the feature set, these abstractions allow low-dimensional visualizations.

We then cluster cell phones based on these abstractions, which allows us to perform (a) cross-category analysis and answer questions such as: *do cell phones with high speed network interfaces have enough computational resources to utilize the network bandwidth? do more flexible cell phones (e.g., that provide support for running third party applications) tend to be more insecure?* and (b) allows prediction of an abstraction (or features) given the knowledge of other abstractions and similar cell phones. This can be useful when complete information about a particular cell phone is not available[6] (e.g., if the cell phone was discontinued several years back but is still being used). For instance, a network service provider may use the network interface of the phone to infer whether the device maybe using an outdated phone, and then use a middlebox to enhance the device security.

5.1 Abstractions for Features

1. **Connectivity (V)**: To characterize the network interface speed offered by a cell phone, we use the average of the maximum data rates[7] offered by the cellular and WiFi interfaces present in a cell phone as follows:

$$V = \frac{1}{2}\left(\frac{R_{cell}}{R_{cell}^{max}} + \frac{R_{wifi}}{R_{wifi}^{max}}\right) \tag{1}$$

   where R_{cell} and R_{wifi} are the maximum data rates supported by the cellular technology and the WiFi interface in a cell phone, respectively. R_{cell}^{max} and R_{wifi}^{max} are the maximum data rates across *all* cellular technologies and WiFi standards found in our dataset, respectively and used for normalization. Observe that a value of V close to zero indicates poor connectivity whereas a value close to 1 indicates that the phone can support very high data rates for both interfaces.

Features	Feature Values
Browser/Markup Language Support	$B_{type} = \begin{cases} 0 & \text{if no browser} \\ 1/4 & \text{if WAP 1.x} \\ 1/2 & \text{if WAP 2.0/XHTML} \\ 3/4 & \text{if HTML supported} \\ 1 & \text{if HTML5 supported} \end{cases}$
Email	$E = \begin{cases} 0 & \text{if no email support} \\ 1 & \text{if email is supported} \end{cases}$
3rd Party Apps	$A = \begin{cases} 0 & \text{if no support for 3rd party apps} \\ 1 & \text{if 3rd party apps are supported} \end{cases}$

Table 5: Categorization of flexibility features.

2. **Capacity (C)**: We characterize the capacity of a cell phone based on its CPU speed and memory size. These features indicate how fast a phone can perform computations as well as the kind of applications it can run. Like with connectivity, we define capacity as the unweighted average of the CPU speed and memory size, which are normalized by their maximum values in the dataset.

3. **Flexibility (F)**: This refers to the ability of a cell phone to run applications like Email, Internet browser, and provide support for running third party applications. We define flexibility, F, as the average of three features:

$$F = (B_{type} + E + A)/3 \tag{2}$$

   where features B_{type}, E, and A and their values are defined in Table 5.

4. **Insecurity (S)**: We define the insecurity of a device as the average of its status (equal to 1 if the device has been discontinued[8] and 0 otherwise) and the ratio of the known vulnerabilities of the device to the total vulnerabilities across all devices in our dataset.

5.2 Clustering

Using the above abstractions for cell phone features, we now cluster phones and carry out a cross-abstraction analysis. Figure 6 shows the visualization of one such clustering using parallel coordinates. Here, we simply cluster the phones based on their connectivity value, using k-means clustering [26] and a cluster size of four[9]. Each parallel vertical axis corresponds to an abstraction and each connected line segments represent a cell phone. We find that for a given level of connectivity, there exists a large range of phone capacities, which indicates that high speed network interfaces are

[4] Other abstractions are indeed possible, however, we choose these abstractions based on the domain knowledge about the features we have seen.

[5] They allow us to reason about a *collection* of related features (e.g., network connectivity as opposed to specific interfaces such as WiFi, 3G, or Bluetooth).

[6] We could not find information about the CPU speeds of ∼70% of the phone models in our dataset even though information about their cellular interface was available.

[7] Using the maximum data rate informs us about the maximum access capacity, which a multi-path transport like MPTCP can utilize under ideal network conditions.

[8] It indicates that the device software would no longer be updated by the vendor and any security vulnerabilities would remain unchecked until the use of the device is discontinued.

[9] We choose the cluster size using the Elbow method [24], which gave a sum of squared error of less than 0.02 for a cluster size of four with only marginal improvement for larger cluster sizes.

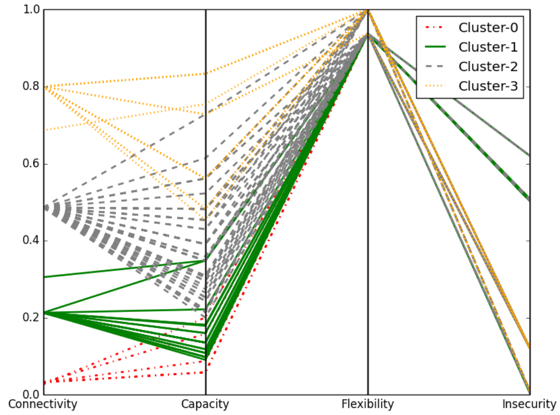

Figure 6: Visualization for connectivity, capacity, flexibility, and insecurity using parallel coordinates. We use the k-means clustering to group abstractions using a cluster size of 4.

not necessarily commensurate with high CPU speeds and/or large memory sizes. We also observe that connectivity value can be a good indicator for predicting flexibility and insecurity but not phone capacity. This observation could be used by ISPs and service providers to infer the security and flexibility characteristics of a device based on its network interface type.

6 Discussion and Implications

How is the distribution of phones likely to change over time? We computed the change in the top-10 phone models (which make up ∼25% of the phones in our dataset) and the top-1000 phone models (making up ∼95% of the phones) across 3 months (i.e., October, November, and December). We found that the change in the top-10 and top-1000 phone models in the last 3 months was 0% and 0.1%, respectively. This suggests that the rate of change in cell phone models is low.

Comparison with other public sources of data and countries: The analysis of our dataset shows that it correlates well with other sources of public data. For example, a report [1] found Android to be the most popular OS for smartphone users in Pakistan. Our dataset showed Android and Symbian to be the most popular OSes among smartphone users. In a remote community in Indonesia, Hartigan et al. [27] found that only ∼16% of the phones were smartphones. While we do not specifically identify smartphones, 12.5% of the users in our dataset had phones with WiFi interfaces, which suggests a similar fraction of smartphones.

Infrastructure support for low-end cell phones: Our analysis reveals insights that can guide infrastructure support needed to improve web experience of users with low end cell phones. One key insight from our work is that changes in the network must be aware of device limitations. For example, a substantial fraction of cell phones are equipped with low data rate cellular interfaces (e.g., GSM, GPRS, and EDGE) and have slow CPUs and small memory sizes. Many such phones provide support for running third party applications. These insights suggest that (a) pre-processing webpages on a proxy server can speed up client-side page load process [35], (b) computations that require high CPU processing and/or memory can be offloaded to the cloud [36], (c) for phones with weak security, providers can provision middleboxes to enhance the device security, and (d) service providers can auto-tune TCP/HTTP parameters based on phone features.

Implications for developing regions research: Our insights reinforce the importance of research on building low cost data communication channels and developing applications suited for low-end phones [20, 25, 29, 30]. For example, Hermes [25] modulates data onto acoustic signals that are sent over a cellular voice call. It provides 1.2 kbps of goodput at a cost per byte that is 50× lower than sending data over SMS, which achieves ∼0.25 kbps of data rate. For cell phones with basic GSM support, such communication channels can provide data service and increase data rates. Similarly, building health, education, and social networking applications for low end phones can benefit a large fraction of users in the developing world. For example, SMSall [16] (previously Chopal) is a mobile social network based on SMS that is reportedly being used by more than 7 million users. Polly [30], a viral telephone-based system aims to reach low-literate population for development-related services through a voice-based game.

Implications for security: Given that a large number of phones (∼58.1%) have been discontinued but are still being used, these phones are vulnerable to a wide range of security attacks. Moreover, as feature information is not available for *all* cell phones, an attacker can use similar analysis of abstractions to correlate similar phones to infer missing feature information. For example, consider two cell phones A and B. Suppose the cellular interface type of both A and B is known. While phone A has been discontinued by the vendor, B's status information is unavailable. An attacker can use the cellular interface type of the phones to infer that B may be outdated. Similarly, by inferring the OS type, an attacker can search through online databases (e.g., [4, 14]) for well-known exploits to launch guided and systematic attacks (e.g., privilege escalation, memory corruption, DoS attacks, and SQL injection).

7 Conclusion

We presented an in-depth study of cell phones from one of the largest cellular providers in Pakistan. Towards this end, we built a custom cell phone database containing 19 cell phone features. We analyzed the security implications of the cell phones being used, identified potential device-level bottlenecks, and discussed how they may impact Internet access in a developing country like Pakistan. To simplify analysis of cell phone features, we proposed abstractions and clustered phones based on these abstractions. Our analysis provides new insights that can inform various stakeholders (e.g., service providers and content providers) in the Internet ecosystem for improving mobile web performance.

8 References

[1] A snapshot of smartphone usage in Pakistan in 2014. https://www.techinasia.com/smartphones-in-pakistan-infographic-2014.

[2] Alexa - Top Sites in Pakistan. http://www.alexa.com/topsites/countries/PK. Accessed: 2016-05-12.

[3] Connectivity Lab. https://info.internet.org/en/story/connectivity-lab/.

[4] CVE Details. http://www.cvedetails.com.

[5] EDGE. http://www.etsi.org/technologies-clusters/technologies/mobile/edge.

[6] Facebook Advertising Research Around the World. https://research.facebook.com/blog/facebook-advertising-research-around-the-world/.

[7] General Packet Radio Service (GPRS). http://www.etsi.org/technologies-clusters/technologies/mobile/gprs.

[8] gsmarena. http://www.gsmarena.com.

[9] HSPA. http://www.3gpp.org/technologies/keywords-acronyms/99-hspa.

[10] IMEI.INFO. http://www.imei.info.

[11] LTE. http://www.3gpp.org/technologies/keywords-acronyms/98-lte.

[12] Mobile technologies GSM. http://www.etsi.org/technologies-clusters/technologies/mobile/gsm.

[13] Mobile technologies GSM. http://www.3gpp.org/technologies/keywords-acronyms/102-gprs-edge.

[14] National Vulnerability Database. https://nvd.nist.gov.

[15] Project Loon. https://www.google.com/loon/.

[16] SMSall: Mobile social network for everyone. http://smsall.pk.

[17] Technologies. https://www.telenor.me/en/About-Telenor/Network/Technologies/.

[18] The State of LTE. https://opensignal.com/reports/2015/09/state-of-lte-q3-2015/.

[19] The World in 2014. ICT Facts and Figures. https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2014-e.pdf.

[20] Anand, A., Pejovic, V., Belding, E. M., and Johnson, D. L. Villagecell: Cost effective cellular connectivity in rural areas. In *ICTD* (2012).

[21] Bischof, Z. S., Rula, J. P., and Bustamante, F. E. In and out of cuba: Characterizing cuba's connectivity. In *IMC* (2015).

[22] Biswas, S., Bicket, J., Wong, E., Musaloiu-E, R., Bhartia, A., and Aguayo, D. Large-scale measurements of wireless network behavior. In *ACM SIGCOMM* (2015).

[23] Danis, C. M., Ellis, J. B., Kellogg, W. A., van Beijma, H., Hoefman, B., Daniels, S. D., and Loggers, J.-W. Mobile phones for health education in the developing world: Sms as a user interface. In *ACM DEV* (2010).

[24] David J. Ketchen, J., and Shook, C. L. The application of cluster analysis in strategic management research: An analysis and critique. *Strategic Management Journal 17*, 6 (1996), 441–458.

[25] Dhananjay, A., Sharma, A., Paik, M., Chen, J., Kuppusamy, T. K., Li, J., and Subramanian, L. Hermes: Data transmission over unknown voice channels. In *MobiCom* (2010).

[26] Hartigan, J. A., and Wong, M. A. Algorithm AS 136: A K-Means Clustering Algorithm. *Applied Statistics 28*, 1 (1979), 100–108.

[27] Heimerl, K., Menon, A., Hasan, S., Ali, K., Brewer, E., and Parikh, T. Analysis of smartphone adoption and usage in a rural community cellular network. In *ICTD* (2015).

[28] Mathur, A., Agarwal, S., and Jaiswal, S. Exploring playback and recording of web-based audio media on low-end feature phones. In *ACM DEV* (2013).

[29] Patra, R., Nedevschi, S., Surana, S., Sheth, A., Subramanian, L., and Brewer, E. Wildnet: Design and implementation of high performance wifi based long distance networks. In *NSDI* (2007).

[30] Raza, A. A., Pervaiz, M., Milo, C., Razaq, S., Alster, G., Sherwani, J., Saif, U., and Rosenfeld, R. Viral entertainment as a vehicle for disseminating speech-based services to low-literate users. In *ICTD* (2012).

[31] Sharma, A., Kaur, M., Koradia, Z., Nishant, R., Pandit, S., Raman, A., and Seth, A. Revisiting the state of cellular data connectivity in india. In *DEV* (2015).

[32] Sharma, A., Kaur, M., Koradia, Z., Nishant, R., Pandit, S., Raman, A., and Seth, A. Revisiting the state of cellular data connectivity in india. In *DEV* (2015).

[33] Velghe, F. Literacy acquisition, informal learning and mobile phones in a south african township. In *ICTD* (2013).

[34] Wang, X. S., Balasubramanian, A., Krishnamurthy, A., and Wetherall, D. Demystifying page load performance with wprof. In *NSDI* (2013).

[35] Wang, X. S., Krishnamurthy, A., and Wetherall, D. Speeding up web page loads with shandian. In *NSDI* (2016).

[36] Wang, X. S., Shen, H., and Wetherall, D. Accelerating the mobile web with selective offloading. In *MCC* (2013).

[37] Zaki, Y., Chen, J., Pötsch, T., Ahmad, T., and Subramanian, L. Dissecting web latency in ghana. In *IMC* (2014).