

A First Look at Performance in Mobile Virtual Network Operators

Fatima Zarinni[†], Ayon Chakraborty[†], Vyas Sekar[°], Samir R. Das[†], Phillipa Gill[†]
[†] Stony Brook University [°] Carnegie Mellon University

ABSTRACT

Recent industry trends suggest a new phenomenon in the mobile market: mobile virtual network operators or MVNOs that operate on top of existing cellular infrastructures. While MVNOs have shown significant growth in the US and elsewhere in the past two years and have been successful in attracting customers, there is anecdotal evidence that users are concerned about cellular performance when choosing MVNOs over traditional cellular operators. In this paper, we present the first systematic measurement study to shed light on this emerging phenomenon. We study the performance of 3 key applications: web access, video streaming and voice, in 2 popular MVNO families (a total of 8 carriers) in the US, where each MVNO family consists of a major base carrier and 3 MVNOs running on top of it. We observe that some MVNOs do indeed exhibit significant performance degradation and that there are key differences between the two MVNO families.

Categories and Subject Descriptors

C.2.1 [Network Architecture and Design]: Wireless Communication; C.2.2 [Network Protocols]: Applications; C.4 [Performance of Systems]: Measurement Techniques.

Keywords

MVNO; Mobile measurements; Mobile performance; Cellular Measurements; Cellular performance; QoE; Applications.

1 Introduction

A new trend has been emerging in the last few years in the cellular market both in the US and in Europe—the rise of mobile virtual network operators or *MVNOs* [12, 14, 15]. At a high-level, MVNOs use the existing cellular infrastructures that are owned by the traditional cellular operators. MVNOs do not incur significant infrastructure or spectrum licensing costs and offer services that are different from traditional cellular operators (e.g., better pre-paid plans and multiple quotas).

While MVNOs started to appear in the market in the early 2000s, they have only recently become more mainstream in terms of market share. The growth is a culmination of several factors: increas-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
IMC'14, November 05–07, 2014, Vancouver, BC, Canada.

Copyright is held by the owner/author(s). Publication rights licensed to ACM
ACM 978-1-4503-3213-2/14/11 ...\$15.00.
<http://dx.doi.org/10.1145/2663716.2663752>.

ing prices of traditional cellular providers for consumers, users' preference in avoiding contractual lock-ins, the popularity of “pre-paid” services, regulatory intervention to ensure competition and market segmentation focusing on niche demographics (e.g., tween markets) [22]. As of Q1 2014, there are 20, 23, 11 and 35 MVNOs running on top of the AT&T, T-Mobile, Verizon and Sprint networks in the US, respectively [5].

Even as MVNOs grow in market share, there are concerns among users about their performance. For example, a quick sampling of popular consumer complaints forums shows significant concerns related to both cost, billing, and service issues (e.g., poor coverage/signal, 3G/4G promised but getting 2G, poor application performance, and frequent disconnections). Shown below are some actual quotes from user forums about MVNOs:

[11]: *I know that AIO is capped at 8mbps download speed. Are all the other MVNOs like Straight Talk, Net10 and Air-Voice also limited to 8mbps download speeds? Do they suffer from higher latency?*

[9]: *I have been throttled every day since last week so each day I lose my 4g/E symbol and once I regain it Im throttled ... I've used 1.4gb and I have only 3 days left on my 30 plan.*

[13]: *Does Sprint have means of degrading service to Ting (and other MVNO) customers in favor of Sprint customers in a particular crowded cell?*

[10]: *My only concern is if the service quality of the service. With Straight Talk, for example, I'd be on AT&T's GSM network in Boston, I think ... but I wonder if as an MVNO customer I'd get second-tier access or service.*

Motivated by the growth of MVNOs and the aforementioned user concerns, this paper presents a first study to shed light on the performance of different MVNOs. While there is a lot of previous work in analyzing mobile performance (e.g., [25, 20, 21, 3]), they have not systematically analyzed performance in MVNOs. To address this gap, we study two major MVNO families in the US. In our study, each family includes the base carrier and three popular MVNOs running on top of the base carrier. While this sample study does not cover all base carriers in US or all MVNOs atop any base carrier, the carrier/MVNO choices have been done systematically, based on popularity (Section 2). In the performance analysis, we hide the actual names of the carriers and MVNOs to protect their business interests. To simplify presentation, we refer to the two base carriers as carrier A and B. We refer to the MVNOs within the carrier A as A1, A2 and A3, and within the carrier B as B1, B2 and B3. The base carrier along with its MVNOs (e.g., A, A1, A2, A3) are referred to as ‘MVNO family’ or just ‘family.’ (e.g., MVNO family A)

As a starting point, we analyze the performance for three dominant usage modes: web access, video streaming and voice calls.

Using over 13,000 measurements collected across 11 locations over a period of 3 months, we address the following questions:

- Does the performance vary across the MVNOs running atop the same base carrier? (e.g., is MetroPCS worse than Straight Talk given that they are both MVNOs running on T-Mobile network?)
- Do MVNOs perform worse compared to the base carrier in each case? (e.g., is H2O Wireless, an MVNO on AT&T network, worse than AT&T?)
- Are there differences across different MVNO families? (e.g., do all MVNOs in a family, say the T-Mobile family, show significantly worse performance than those in another family, say the AT&T family?)

We analyze application-specific quality-of-experience (QoE) metrics to address these questions. We also perform factor analysis to correlate the observed application-level performance with network-level performance, such as, TCP throughput, round-trip times (RTTs), packet loss rates, DNS look up times, and PHY-layer characteristics to attribute the observed performance differences (if any) to structural differences across the operators.

Our key findings are:

- The base carrier often performs better than the MVNOs and sometimes significantly so. For instance, some MVNOs over base carrier B fail to load a non-trivial ($\geq 10\%$) fraction of YouTube video requests and can have up to $6\times$ worse page load time.
- There is significant diversity across MVNOs within the same MVNO family, for both the A and B MVNO families. For instance, often B2 performs considerably worse than B1 and B3 in MVNO family B.
- There are non-trivial differences between the two MVNO families; overall the MVNOs running atop A have better performance w.r.t the base carrier compared to their B counterparts.
- Finally, we see key differences across applications as well. While voice quality is largely similar across all MVNOs and base carriers, there is huge discrepancy in data performance both for web access as well as video streaming.

We hope that this paper serves as a motivation for future large-scale measurement studies in this direction, that would span wider areas, larger number of MVNOs and wider variety of data plans.

2 Measurement Setup

In this section, we begin by describing the choice of phone, carriers, and cellular plans. Then, we describe our data collection methodology.

Choice of phone: To ensure we do not have phone-specific effects (e.g. CPU speed/memory access latency/cache size) in our measurements, we use the same phone model for all carriers – Google’s Nexus 4 with 2GB RAM, Quad-core CPU and 2G/3G/4G support (i.e., EDGE/UMTS/HSPA/HSPA+). All of our phones run the Android 4.2.2 (JellyBean) OS. Since Nexus 4 only supports GSM-based carriers, this study is limited to such carriers and their MVNOs only. We leave the investigation of performance in CDMA-based carriers and their MVNOs for future work.

Choice of carriers and plans: We chose popular and widely-used MVNOs that run atop two major base carriers in the U.S. We call them carriers A and B, respectively. We used Google Trends [4] and the list of all the available MVNOs [5] to find the top 3 MVNOs for each of these base carriers. The 6 MVNOs are summarized in Table 1. A1, A2 and A3 run atop A; B1, B2 and B3 run atop

| Carrier | Type | Plan (all pre-paid except B) | \$/Mo. |
|---------|------|--------------------------------------|--------|
| A | Base | Unlimited talk/text, 2.5GB data @ 4G | 60 |
| A1 | MVNO | Unlimited talk/text, 2.5GB data @ 4G | 45 |
| A2 | MVNO | Unlimited talk/text, 3GB data @ 4G | 50 |
| A3 | MVNO | Unlimited talk/text, 2.5GB data @ 4G | 60 |
| B | Base | Unlimited talk/text, 2GB data @ 4G | 65 |
| B1 | MVNO | Unlimited talk/text, 2.5GB data @ 4G | 50 |
| B2 | MVNO | Unlimited talk/text, 2GB data @ 4G | 50 |
| B3 | MVNO | Unlimited talk/text, 2GB data @ 4G | 50 |

Table 1: Mobile carriers and plans used in our study.



Figure 1: We conduct measurements at 11 different locations spanning across a 3000 km² wide area. The annotations show the names of the measurement locations along with the type of location, and the number of measurement sets.

B. Cellular providers offer a range of plans with different prices. Hence, to provide a fair comparison between carriers, we select *similar* plans for all the carriers (as summarized in Table 1), in terms of features. When the exact plan was not available we picked the closest comparable plan.

Data collection: We selected 11 reasonably diverse locations spanning different usage scenarios: urban/suburban, shopping areas, residential, office/lab and hospital. Figure 1 shows the geographical spread of our measurement locations. We acknowledge a potential limitation, that all our measurements occurred in the Long Island/New York region. However, this region is a major population hotspot, covering part of New York city metro area and associated suburbs.

We developed a suite of custom scripts and mobile applications for web browsing, video streaming, and voice calls, representing common usage modes. Our custom tools collect relevant user quality-of-experience (QoE) metrics, and we defer application-specific details to the following sections.

At each location, we use four identically configured Nexus 4 phones (one for base and three for MVNO carriers) to run the same suite of experiments concurrently at that location. Our scripts run these applications at each location typically hourly or half-hourly for most of the day – often starting at early morning and going until late night – over different days of the week modulo practical constraints (e.g., shop/mall closures).

On average, we conducted about 150 sets of measurements for each carrier, across different locations, during Jan-Mar 2014, on different days. Each measurement set consists of a series of application runs, e.g., web page access for a set of chosen web sites, video streaming, voice calls, TCP upload throughput test, etc.

Concurrent with the QoE measurements, we also log packet traces using the `tcpdump` tool and a range of relevant phone characteristics using the Android API (e.g., radio stats), to enable further factor analysis. We verified separately via the `top` utility, that this additional monitoring adds only a modest CPU overhead ($\approx 5\%$). This does not bias our measurements. Prior to conducting actual measurements, we performed tests over WiFi, where we ran our apps with and without additional logging, and we measured the performance for web, video and voice applications, as well as, network tests. The attained results showed negligible difference in performance with this logging enabled or disabled.

Our analysis did not reveal any significant location, time-of-day or day-of-week specific change in terms of performance of carriers with respect to each other. Thus, we present only aggregate statistics (over all locations, times and days) and focus on performance differences across carriers and MVNO families. Since the experiments for base and MVNO carriers are always colocated in both space and time, we believe it is a fair characterization of the performance issues we describe in the rest of the paper.

3 Application Performance

In this section, we analyze the performance of the MVNOs and the base carriers for three common modes of mobile usage: web access, video streaming, and voice calls. In each case, we describe the application-specific setup, and the relevant QoE metrics we measure. We also correlate the observations to key network-level and PHY-layer characteristics.


3.1 Web Browsing

Setup and QoE Metric: We choose six popular websites with diverse characteristics: Youtube, Amazon, Wikipedia, Twitter, Bing, and CNN. All of these sites had an overall Alexa rank ≤ 20 in April 2014. We developed a custom browser application using Android WebViewClient. At each measurement site, the app visits each website’s mobile landing page (in random order across carriers) and records the *page load time* QoE metric. We measure page load time as the difference between the time the URL is requested from the browser and the time when all the web objects (html text, images, etc.) are fetched and the `onPageFinished` event [7] is triggered.


Note that the set of webpages accessed is diverse in terms of structure and content size, with CNN and Amazon constituting the two largest median content sizes in the set (≈ 570 KB and 400 KB, respectively) and Twitter and Bing having the smallest (≈ 89 KB and 100 KB, respectively).

Evaluation of Page Load Times: Figure 2 shows the distribution of page load time across all runs for the six websites.¹ There are three key observations here. First, typically the carriers in MVNO family A perform better than their B counterparts; e.g., for CNN all carriers in MVNO family A perform better than all carriers in MVNO family B, and sometimes significantly so. Second, while the differences between base carrier A and its MVNOs are only modest, we see significant differences between base carrier B and some of its MVNOs; e.g., B2 is almost 6 \times worse than B for CNN. Finally, we see non-trivial variability across MVNOs within the same MVNO family; e.g., B2 is often considerably worse than other MVNOs in MVNO family B, and A1 is slightly worse than other MVNOs in MVNO family A. We confirmed that these differences between carrier page load times are statistically significant

¹Note that, >10 sec page load times are not surprising on a mobile platform, as seen in prior work [38, 19, 36]. For example, Welsh reports 75 seconds page load time for a webpage over a cellular link [38].



(a) MVNO family A



(b) MVNO family B

Figure 2: Distribution of page load times (median, 25th and 75th percentiles): We see that (a) MVNO family A usually performs better than MVNO family B; (b) within each MVNO family one or more MVNOs is worse than the base carrier; and (c) some MVNOs (e.g., B2, B3) suffer more than others).

using the Kolmogrov-Smirnov (K-S) [8] statistical test, but do not show these results for brevity.

Factor Analysis: To understand the causes of these performance differences, we looked at lower-layer metrics such as DNS lookup time, RTT, TCP retransmission rates, and signal strength. We computed the Pearson’s correlation between the difference of page load times for the base carrier and the MVNOs and that of different lower-layer metrics, for every website. Based on this analysis, we zoom in on two key factors – RTT and TCP retransmissions (Figure 3). First, we can see that MVNO family A has generally lower RTTs and retransmission rates than MVNO family B. As prior studies have shown, lower RTTs imply lower page load times, which is consistent with our observations that A and its MVNOs have lower page load times [18]. Second, we see in Figure 3a that within the MVNO family A, A1 which had higher page load times, indeed has higher RTTs.² Finally, Figure 3d shows that the MVNOs in MVNO family B (B2 and B3) with the highest page load times see very high retransmission rates.

We also observe that B1 has the lowest retransmission rates in its MVNO family, however, still higher RTTs than B, thus resulting in B1 having a lower page load time than B2 and B3, but, higher than B. However, this still does not explain some of the very high (> 30 s) page load times for B1 (e.g., CNN in Figure 2b). Further analysis of the packet traces showed significant TCP idle times as shown in one example timeseries in Figure 4a. Figure 4b breaks down the page load measurements in the B MVNO family in two

²Higher RTTs for both Twitter and Bing, as compared with other webpages, could be due to the content-server locations that were accessed for these sites, or due to the path from carrier A’s gateway routers to these servers [35, 39].

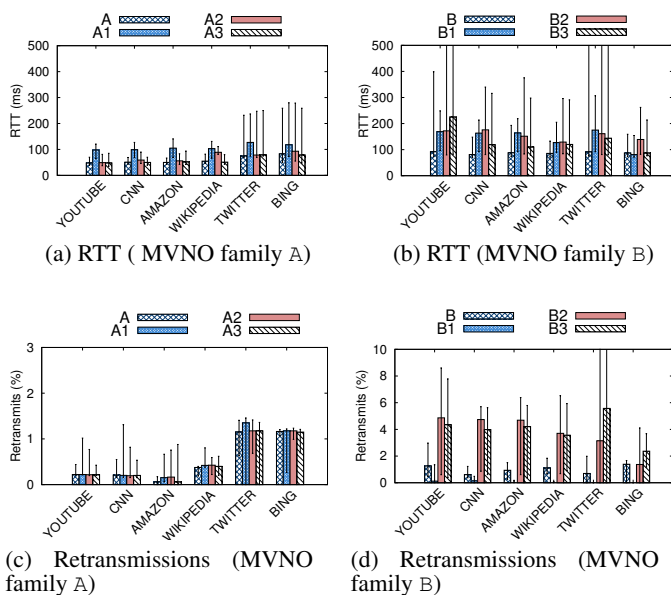


Figure 3: Focusing on the key observed factors shows that generally speaking the MVNOs in family A with higher page load times have higher RTT and the MVNOs in family B with higher page load times tend to have high retransmission rates.

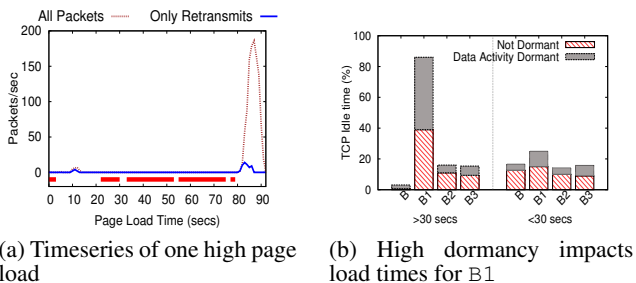


Figure 4: Higher page load times for B1 relative to B are not due to retransmissions but rather due to high radio dormancy periods. The red line in (a) shows intervals when data activity is dormant.

bins (< 30 s and > 30 s) and shows that these TCP idle periods have non-trivial influence on the page load times. This is specifically true for B1 where the long page load times have about 80% idle periods.

Further inspection reveals that many (but not all) of these idle periods are actually due to physical link being dormant, (as revealed by the `DATA_ACTIVITY_DORMANT` flag [1]). We suspect that this is influenced by the RRC state machine at the radio layer as defined in the 3GPP standard [24], but we do not have visibility to actual RRC states using the commodity Nexus 4 phone to examine this further. Prior work (e.g., [19, 34]) has also shown that inappropriately tuned RRC states impact web access performance. Overall, this suggests some potential misconfiguration or service differentiation at the radio layer for the MVNO B1 running over carrier B. In contrast, TCP idle/dormancy issues are negligible for MVNO family A and are not shown.

We also analyzed signal strengths, handoffs and the pool of cell-ids that the carriers are associating with and found no significant differences between carriers within the same MVNO family. This implies that these radio-layer aspects did not play a significant role in the performance difference observed between MVNOs. Some

prior work (e.g., [23, 29]) also noticed little correlation between signal strengths and performance when analyzing their collected measurements. This is likely due to the signal strengths usually falling above a certain threshold.

Our investigation in this section also revealed interesting information about the structural differences across the A and B MVNO families. In MVNO family B, all web traffic goes through an explicit proxy server that terminates TCP connections while MVNO family A appears to use a transparent proxy that relays the connections to the webservers.³

3.2 Video Streaming

Setup: We choose a 3-minute YouTube video available in both high/low quality and play it in a custom app. We use the Android YouTube APIs [2] to extract player states (paused, playing, buffering) to compute the QoE metrics described below. Similar to the web experiments, we run measurements for both the base carrier and associated MVNOs, simultaneously, at multiple locations and at different times of the day.

QoE metrics and Evaluation: The key video QoE metrics are: (1) *video resolution* being delivered;⁴ (2) *startup delay* or the time between the user clicking on the play button and the time the video starts playing; (3) *buffering ratio* or the percentage of the session duration spent in buffering state; and (4) *load failures*, where the video fails to load. Figure 5 summarizes the distribution of these metrics.

First, we observe that, with respect to resolution quality, carrier B and its MVNOs always use the high-resolution version of the video. On the other hand, carrier A and its MVNOs play a mix of resolutions, except for A3 that always plays the lower resolution video. Second, in terms of startup delay, MVNO family B overall shows a higher startup delay than MVNO family A, for the higher-resolution cases. Also, consistent with the web measurements, we find that in MVNO family B, the base carrier outperforms its MVNOs in terms of buffering ratio as well, and B2 again performs the worst amongst the MVNOs in its family. Finally, we find a non-trivial number of video load failures for the MVNOs in the B family; e.g., B1 fails $\approx 20\%$ of the time.

Factor Analysis: As before, we use the correlation coefficients to zoom in on key network-level factors. The startup delay and buffering states are (unsurprisingly) mostly influenced by *TCP throughput*. Figure 6 shows the difference in the measured TCP throughput across the carriers and confirms the earlier observations about video QoE. Surprisingly, MVNO family B chooses the higher quality video even though it has lower TCP throughputs than MVNO family A (and hence incurs more buffering). We suspect that this is related to the explicit proxy described earlier; i.e., the bitrate negotiation at the beginning of the session is done by the proxy and does not account for the actual “last hop” throughput achievable by the client.

To further understand the load failures, we analyzed the packet-level traces and find two reasons behind these failures: (1) the proxy *blocks* the video requested by the client by sending an HTTP

³We were able to detect the transparent proxy using Netalyzr which showed HTTP header modifications [6].

⁴The YouTube API does not perform dynamic video resolution adaption on mobile. It selects a resolution that it considers suitable for the current connection at the start and uses it for the entire session.




Figure 5: Video quality-of-experience metrics for the MVNOs and base carriers. Note that MVNO family B always plays the high-quality resolution and suffer significant buffering, startup delay, and video load failures.




Figure 6: TCP throughput influences video quality

response with the status code 403 (‘access forbidden’) and (2) the proxy does not respond to the initial request from the client causing the client to timeout. Specifically, B1 experiences the largest number of type (2) video load failures; this is related to radio dormancy issues discussed in Section 3.1.

3.3 Voice calls

Setup and QoE metrics: We created a custom *auto dial* application that is scripted to automatically call a number. This application runs on all 4 phones at different locations and times of the day. We also setup a recipient phone in the lab, and we build and run another custom application on this phone to log the time of the first ring, accept the call, and then immediately end the call. Because the Android top-level framework does not allow us to automatically answer/end a call, we mimicked a Bluetooth Headset request to automatically answer the call and used lower-layer APIs for ending the call. With this setup we compute the *Call Setup Time* or the time elapsed between the time the caller makes a call and when the callee receives it. To ensure that the caller/receiver are in sync, we use the `ClockSync` [16] Android app. We separately verified that the synchronization error was ≤ 10 ms (not shown); this suffices for our analysis as we look for user-perceivable (e.g., ≥ 100 ms) differences in performance.

To measure the *audio quality*, we establish calls between each of the 4 phones and a Google Voice number on a laptop. We play a 3 minute (based on average audio call durations) audio file on the laptop and record the incoming audio to the phone. To minimize background noise we direct the audio output from the phone to the recorder via a standard 3.5mm cable. We compute the cross-correlation of the Mel-frequency Cepstral Coefficients (MFCCs), (recommended in the audio/speech processing literature [27]), between the reference audio file and the recorded audio file. We normalize this value by dividing it by the score attained when cross-correlating the MFCCs of the original file with itself, and we call this normalized value the *Call Quality Score*.

Evaluation: Figure 7a shows the distribution of the call setup time for the different carriers. MVNO family A showed fairly similar




Figure 7: Call quality in terms of setup time and the audio quality. While there is no significant difference in the call quality, we do observe that some of the MVNOs in MVNO family B have a higher call setup time.

values for call setup times (median of 5-6 seconds). With MVNO family B, we notice that B2 and B3 have a 1.5 second higher median call setup time. However, this difference cannot be attributed to any client-side metric we collected. Figure 7b shows the Call Quality Score for the two MVNO families. In this case, we do not observe significant differences across the providers.⁵ Since the discrepancy in quality is low, unlike the data experiments, we do not perform any further factor analysis.

4 Other Applications

In addition to the three usage scenarios discussed in the previous section, we also conduct smaller-scale measurements to capture other common user concerns. We briefly summarize the main observations from these experiments.




Figure 8: TCP Uplink Throughputs.

Uploads: Web and video workloads are largely download-bound. Users are increasingly using phones to upload content (e.g., Instagram, Vine). To understand the impact on such applications, we

⁵Note that Nexus 4 phones do not support VoLTE (Voice over LTE). Hence, audio voice calls are sent over a channel separate from the data channel, and thus voice is not impacted by differences on the data channel.

measure the upload speeds obtained by different carriers to a reference campus server in Stony Brook. Figure 8 shows similar characteristics to our previous experiments—MVNOs in MVNO family B perform significantly worse than base carrier B, whereas base carrier A and its MVNOs perform roughly similarly.

Video Chat (Google Hangout): We pick a popular video chat application – Google Hangout and evaluate its performance across carriers. We establish a chat (5 minutes long) from the phone to a well-provisioned laptop and play a video in front of the phones and the laptop using another screen. We collect packet traces at both ends. We repeat our experiments at 3 different locations, lab and two residential areas. We analyze frames/sec received at the laptop as well as the sending and receiving UDP throughputs. We did not observe any significant difference between the base carrier and their MVNOs for both A and B MVNO families. One interesting observation is that video chat shows no performance differences on MVNOs in MVNO family B, which is unlike the case of video streaming (Section 3.2). We speculate that this is due to a combination of two reasons: (1) chat traffic uses UDP and does not go through the explicit proxy and (2) unlike the YouTube video which chooses a static bitrate at the start, Hangout uses dynamic bitrate adaptation. (Recall that most of the problems in MVNO family B was the poor choice of initial bitrate via the proxy.)

Traffic Shaping and Port Blocking: Users like to know if MVNOs throttle or block less common applications. This is particularly relevant as we have seen use of proxies (Section 3.1) and use of middleboxes in cellular networks [37] is well-known. We use two tools, Bonafide [17] and Netalyzr [6], as they provide complementary coverage over the set of application tests. We used this to study 3 different types of applications: (1) BitTorrent, (2) VoIP-H323 and (3) RTSP-based apps. We found no evidence of application-specific traffic shaping for these protocols, in both MVNO families. Additionally, Netalyzr reveals that while MVNO family A does not exhibit any port blocking, MVNO family B exhibits more diverse blocking behavior. For example, B blocks TCP-based SIP and UDP access to NetBIOS-NS servers. While B3 and B2 do not block any ports, B1 blocks many application ports (FTP, PPTP, NTP, NetBIOS-NS, NetBIOS-DGM, IKE Key Exchange).

Coverage: As seen in our user quotes, users want to know if MVNOs get the same coverage/treatment as the base carriers (e.g., [10]). As discussed earlier, we logged relevant lower-layer information—serving cell-id, signal strength (RSSI/RSCP/RSRP), link layer technology used (e.g., EDGE, HSPA, HSPA+). In addition, we did a number of driving experiments covering major routes within the map in Figure 1. We found that in general, the carriers in each MVNO family connect to a similar set of cell-ids in a given location and that there was no statistically significant difference in signal strength or link-layer technology used.

Quota usage: Another common concern for users is whether carriers start throttling before the actual usage quotas are reached (e.g., [9]). We correlated the performance for different applications vs. the data usage amount for every billing cycle. We did not observe throttling behavior for either MVNO family. A detailed study of this subject via more controlled experiments is an interesting direction of future work, especially in light of known accounting discrepancies (e.g., [32]).

5 Related Work

With the growth of mobile traffic, there are several prior and ongoing efforts in mobile measurement. While the tools and techniques

they use are similar to our work, the key difference is that these have not focused on the MVNO phenomenon to characterize differences across MVNOs or MVNOs vs. base carriers.

Mobile measurements: Previous studies have measured mobile performance from the infrastructure-side [26, 28, 31, 33] and the client-side [20, 21]. These focus primarily on characterizing traffic usage patterns, which is orthogonal to our work. Wang et al. showed how middlebox effects (e.g., timing out idle TCP connections) can have a huge impact on the mobile application performance [37]. Huang et al. compare different carriers on a range of applications across different smartphone hardware [25]. However, their study did not cover MVNOs. More recent studies analyze performance variability within carriers [30] and diagnose causes of high latency in cellular networks [39]. These are interesting factors to further dissect MVNO performance.

Tools and datasets: Several crowd-sourced solutions gather mobile measurements; e.g., FCC’s broadband measurement tool [3], OpenSignal (www.opensignal.com), Mobiperf (www.mobiperf.com), OOKLA Speed Test (www.speedtest.net). These focus mostly on network-level metrics (e.g., latency, throughput, signal strength) and do not measure user-perceived QoE metrics which is our primary focus. Bashko et al. developed the Bonafide tool to detect traffic shaping and service differentiation [17]. Netalyzr is also a powerful tool for detecting port blocking, proxies, and DNS issues [6]. We leverage these two tools and apply them to study MVNOs.

6 Conclusions

In this paper, we presented a first attempt to shed light on a recent and growing trend in the mobile market: mobile virtual network operators or MVNOs. While these have been growing in market share, there are natural concerns about their performance and there has been little work done on systematically understanding this area. To fill this gap, we conducted a systematic measurement study with two major MVNO families in the US. Our analysis shows that while the MVNOs share the network infrastructure of the base carriers, there is visible performance degradation in quality of experience metrics for common mobile phone applications for some MVNOs. Further, MVNOs in the same MVNO family do not perform equally, and the two MVNO families behave differently. Deeper analysis reveals a range of structural and lower-layer differences across MVNO families and MVNOs, including use of proxy, varying latencies and loss rates, data activity dormancy issues and various forms of blocking/denials. We hope that our observations motivate and trigger future deeper and large-scale studies, across larger regions, more MVNOs and more variety of data plans, perhaps by using mobile measurement platforms being deployed in the wild.

7 Acknowledgements

We thank Gaurav Dugar, Su’aad Zaman, Naila Kabiraj, and Jihoon Ryoo for their help with data collection. We also thank Bashko et al., for sharing their BonaFide tool [17] with us. This research was partially supported by NSF grant CNS-1117719.

8 References

- [1] Android TelephonyManager Class.
http://developer.android.com/reference/android/telephony/TelephonyManager.html#DATA_ACTIVITY_DORMANT.

- [2] Android YouTube Player API. <https://developers.google.com/youtube/android/player/>.
- [3] FCC - Measuring Broadband America. <http://www.fcc.gov/measuring-broadband-america>.
- [4] Google Trends. <http://www.google.com/trends/>.
- [5] List of MVNOs in the U.S. http://en.wikipedia.org/wiki/List_of_United_States_mobile_virtual_network_operators.
- [6] Netalyzer, HTTP proxy detection through HTTP header or content modification. http://nl.netalyzr.icsi.berkeley.edu/info_httpproxy_header.html.
- [7] onPageFinished method: Android WebViewClient. <http://bit.ly/1l0eaVP>.
- [8] The K-S Statistical Test. http://en.wikipedia.org/wiki/Kolmogorov-Smirnov_test.
- [9] Straight Talk has upped the ante on throttling users! <http://www.howardforums.com/showthread.php/1773165-Straight-Talk-has-upped-the-ante-on-throttling-users!>, 2012.
- [10] Is there a catch to MVNOs? <http://ask.metafilter.com/235587/Is-there-a-catch-to-MVNOs>, 2013.
- [11] AT&T vs. other MVNO's on latency, capped download speeds . <http://www.howardforums.com/showthread.php/1813979-AT-amp-T-vs-other-mvno-s-on-latency-capped-download-speeds>, 2013.
- [12] M2M and MVNOs driving US connections growth. <https://gsmaintelligence.com/analysis/2013/08/m2m-and-mvnos-driving-us-connections-growth/397/>, 2013.
- [13] Roaming, Regional, Mobile to Mobile, and Other Questions. <https://help.ting.com/entries/21227328-roaming-regional-mobile-to-mobile-and-other-questions>, 2014.
- [14] Sprint, T-Mobile execs explain the MVNO explosion. <http://www.fiercewireless.com/special-reports/mvno-explosion-will-latest-wave-virtual-operators-survive-0>, 2014.
- [15] Western Europe-Evolution of MVNO Business Models-Market Growth, Trends and Opportunities. <http://www.releasewire.com/press-releases/western-europe-evolution-of-mvno-business-models-market-growth-trends-and-opportunities-493894.htm>, 2014.
- [16] S. Baranov. Android ClockSync Application. <https://play.google.com/store/apps/details?id=ru.org.amip.ClockSync&hl=en>.
- [17] V. Bashko, N. Melnikov, A. Sehgal, and J. Schonwalder. Bonafide: A traffic shaping detection tool for mobile networks. In *IFIP/IEEE IM*, 2013.
- [18] M. Belshe. More Bandwidth Doesn't Matter (much). <http://bit.ly/1rNC01C>, 2010.
- [19] J. Erman, V. Gopalakrishnan, R. Jana, and K. Ramakrishnan. Towards a SPDY'ier mobile web ? In *ACM CoNext*, 2013.
- [20] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin. Diversity in smartphone usage. In *ACM MobiSys*, 2010.
- [21] H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin. A first look at traffic on smartphones. In *IMC*, 2010.
- [22] K. Fitchard. Why are MVNOs so hot right now? Thank the carriers. <http://gigaom.com/2012/06/25/why-are-mvnos-so-hot-right-now-thank-the-carriers/>, 2012.
- [23] A. Gember, A. Akella, J. Pang, A. Varshavsky, and R. Caceres. Obtaining in-context measurements of cellular network performance. In *Proceedings of IMC 2012*.
- [24] I. Grigorik. *High Performance Browser Networking*. Chapter 7, O'Riley, 2013.
- [25] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl. Anatomizing application performance differences on smartphones. In *ACM Mobisys*, 2010.
- [26] Y. Lee. Measured TCP Performance in CDMA 1x EV-DO Network. In *PAM*, 2006.
- [27] M. Lindasalwa, M. Begam, and I. Elamvazuthi. Voice recognition algorithms using Mel Frequency Cepstral Coefficient (MFCC) and Dynamic Time Warping (DTW) techniques. In *arXiv preprint arXiv:1003.4083*, 2010.
- [28] G. Maier, F. Schneider, and A. Feldmann. A first look at mobile hand-held device traffic. In *PAM*, 2010.
- [29] J. Manweiler, S. Agarwal, M. Zhang, R. Roy Choudhury, and P. Bahl. Switchboard: A matchmaking system for multiplayer mobile games. In *Proceedings of ACM Mobisys*, 2011.
- [30] A. Nikraves, D. Choffnes, E. Katz-Bassett, Z. M. Mao, and M. Welsh. Mobile network performance from user devices: A longitudinal, multidimensional analysis. In *PAM*, 2014.
- [31] U. Paul, A. Prabhu Subramanian, M. Buddhikot, and S. Das. Understanding traffic dynamics in cellular data networks. In *INFOCOM*, 2011.
- [32] C. Peng, C. yu Li, G. hua Tu, S. Lu, and L. Zhang. Mobile Data Charging: New Attacks and Countermeasures . In *Proc. CCS*, 2012.
- [33] X. Qiang, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman. Identifying diverse usage behaviors of smartphone apps. In *IMC*, 2011.
- [34] S. Rosen, H. Luo, Q. A. Chen, Z. M. Mao, J. Hui, A. Drake, and K. Lau. Discovering Fine-grained RRC State Dynamics and Performance Impacts in Cellular Networks. In *Proceedings of ACM Mobicom*, 2014.
- [35] J. P. Rula and F. E. Bustamante. Behind the curtain: The importance of replica selection in next generation cellular networks. In *Proceedings of ACM SIGCOMM*, 2014.
- [36] A. Sivakumar, V. Gopalakrishnan, S. Lee, S. Rao, S. Sen, and O. Spatscheck. Cloud is Not a Silver Bullet: A Case Study of Cloud-based Mobile Browsing. In *Proceedings of ACM HotMobile*, 2014.
- [37] Z. Wang, Z. Qian, Q. Xu, Z. Mao, and M. Zhang. An untold story of middleboxes in cellular networks. In *ACM SIGCOMM*, 2011.
- [38] M. Welsh. Why Mobile Performance is Hard. ACM Mobicom Industrial Paper Session . http://www.sigmobile.org/mobicom/2013/MobiCom2013_IndustrialTalks_MattWelsh.pdf, 2013.
- [39] K. Zarifis, T. Flach, S. Nori, D. Choffnes, R. Govindan, E. Katz-Bassett, Z. Mao, and M. Welsh. Diagnosing path inflation of mobile client traffic. In *PAM*, 2014.