GENI WiMAX Performance: Evaluation and Comparison of Two Campus Testbeds

Fraida Fund*, Cong Wang†, Thanasis Korakis*, Michael Zink†, Shivendra Panwar*
*Department of Electrical and Computer Engineering
Polytechnic Institute of NYU
ffund01@students.poly.edu, korakis@poly.edu, panwar@catt.poly.edu
†Department of Electrical and Computer Engineering
University of Massachusetts Amherst
{cwang, zink}@ecs.umass.edu

***Abstract*—**

In the last few years, there has been an increasing awareness of the need to evaluate new mobile applications and protocols in realistic wireless settings, and platforms such as the GENI WiMAX testbeds have been developed to fulfill this need. However, wireless testbed users have experienced frustration when straightforward usage scenarios do not consistently agree with the high data rates that are advertised by the wireless technology. This work seeks to clarify the performance characteristics of two GENI WiMAX testbeds under various wireless signal conditions and network traffic patterns. By measuring the performance of several popular wireless Internet applications in two very different wireless environments, we gain a deeper understanding of how a researcher may expect the GENI WiMAX platform to behave. Our findings include some counterintuitive results, e.g. that increasing signal quality can reduce application throughput, and that applications using a single TCP flow may achieve as much as 72% less throughput than an application in an identical setting that uses multiple TCP flows. With this work, we hope to help other researchers design realistic experiments on wireless Internet systems, understand the perceived shortcomings of the GENI WiMAX platform, and interpret their experimental results in the context of the wireless setting in which the experiment was conducted.

I. INTRODUCTION

The astonishing growth of mobile computing and the proliferation of mobile devices has created an increasing need for research platforms composed of real wireless networks and devices. Because of the inherent difficulty in conducting experiments with real mobile devices and wireless networks (especially cellular networks), a great deal of published work to date is based on results from simulations, which necessarily ignore much of the complexity inherent in a realistic wireless network and therefore produce results that are not always reliable or reproducible [1] [2]. Fortunately, as mobile computing becomes more pervasive, there has been an increasing awareness of the need to evaluate new applications and protocols in more realistic settings, and various platforms have been developed to meet this demand [3] [4].

The GENI [5] WiMAX initiative aims to fulfill this need with the creation of an open, programmable, mesoscale cellular infrastructure compatible with the IEEE 802.16e ("mobile WiMAX") standard [6]. This technology offers wide-area broadband wireless communication similar to other 4G systems, with advanced mobility and flexibility, high data rates, and support for diversified QoS service classes. The bandwidth and range of WiMAX make it suitable for a number of potential applications, including providing mobile broadband coverage to a variety of devices as a Metropolitan Area Network (MAN), providing an alternative to cable for last mile broadband access (especially to rural areas where it is not practical to install a wired connection), and serving as a wireless backhaul for 2G, 3G, and 4G cellular networks.

Despite the promise of WiMAX and other 4G wireless systems, WiMAX testbed users have experienced frustration when their straightforward usage scenarios do not consistently meet the high data rates that have been reported for these technologies. Perceptions about what a modern wireless system provides as a network platform are often oversimplified, even among networking researchers. In particular, expectations about high download rates are likely to be satisfied only under certain conditions. There has been no work to date that clarifies how a researcher may expect the GENI WiMAX platform to behave, given the behavior of the application or protocol to be evaluated and the wireless context of the experiment.

The Polytechnic Institute of NYU (NYU-Poly) and the University of Massachusetts Amherst (UMass Amherst) were among the first group of campuses to deploy the GENI mesoscale testbed. A systematic study of WiMAX at these two sites offers unique insights on the performance of the platform from the perspective of application performance, by measuring the performance of several popular wireless Internet applications on this platform, and from the perspective of wireless behavior, by repeating this evaluation under diverse wireless environments. The measurements presented in this work come from several classic wireless Internet usage scenarios, including a traditional server-client file transfer, a peer-to-peer file transfer, and a video streaming scenario. The experiments were conducted in two very different wireless propagation settings - the dense urban area of NYU-Poly and the semi-rural UMass Amherst campus.

The analysis of these measurements seeks to address the following questions:

- What characteristics suggest that an application will perform well in a wireless setting, and what characteristics suggest it will perform poorly?

- What is the relationship between the performance of an application and wireless signal quality? Does increased signal quality always improve application performance?
- How much of an effect does the wireless propagation environment have on application performance?

The rest of this paper is organized as follows. The WiMAX infrastructure in place at each of the campuses is described in Section II. Section III gives background information on the capacity of WiMAX systems under different conditions. Experimental methodology and results are given for experiments on fixed WiMAX nodes in Section IV, and for mobile WiMAX nodes at both campuses in Section V. Section VI places this work in the context of related work. Finally, Section VII concludes with our direction for future work.

II. Overview of GENI WiMAX Testbed and GIMI Measurement Infrastructure

This work was conducted on GENI WiMAX testbed facilities, using software tools from the WiMAX project as well as the GIMI [7] instrumentation and measurement infrastructure. A brief overview of the WiMAX testbed infrastructure and the relevant components of the GIMI toolset is given here.

The WiMAX testbeds at both campuses include a WiMAX base station (BS) and a number of WiMAX-equipped clients (nodes). The testbeds are managed by the cOntrol and Management Framework (OMF) [8], which gives local and remote researchers a robust set of scripting, experiment control, management, and measurement tools to run their experiments. The testbed nodes may be “fixed” nodes, which remain connected to a wired control network at all times, or “mobile” nodes, which may be temporarily disconnected from the control network for mobile experiments.

The GENI WiMAX BSs used in this study are commercial WiMAX 802.16e radios from NEC, operated via customized control software developed at WINLAB [9] that is installed on a separate BS controller. The BSs at NYU-Poly and at UMass Amherst both operate in licensed frequency centered at 2.595 Ghz. Table I summarizes key BS configuration parameters for the experiments presented in this paper.

A. Testbed Configuration

Access Mode	SOFDMA/TDD
Center Frequency	2.595 GHz
Channel Bandwidth	10 MHz
Transmit Power	38 dBm
TDD DL:UL ratio	35:12 (symbols)
CRC	Enabled
Packing	Enabled
Fragmentation	Enabled
Compressed MAP	Disabled
Service Class	Best Effort (BE)

TABLE I: Selected parameters of WiMAX BS configuration.

The network configuration for the WiMAX testbeds is shown in Fig. 1. Each testbed node is equipped with an Intel Centrino Advanced-N+WiMAX 6250 wireless network adapter which is configured to connect to a GENI-operated WiMAX network. When a testbed node requests entry to the WiMAX network, Generic Routing Encapsulation (GRE) [10] tunnels are set up on the BS and the BS controller to route traffic to and from the WiMAX client. Meanwhile, a Click software router [11] running on the BS controller is configured to forward traffic between the WiMAX network and a user-defined datapath. In this way, WiMAX clients can communicate with other hosts on a campus network, the public Internet, or a GENI backbone network.

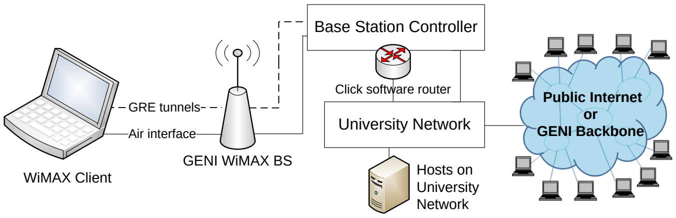

Fig. 1: The WiMAX testbed network configuration. A software router configured by the BS controller forwards traffic from each client to its predefined datapath on the university network, the public Internet, or a GENI backbone link.

B. Measurement Infrastructure

This work was made possible by the extended measurement capabilities offered by several open-source tools. The GIMI instrumentation and measurement toolset provided storage and presentation of measurement data, which was collected by specially instrumented versions of popular Internet applications.

1) GIMI measurement data archive and presentation utilities: The GIMI project [7] aims to provide instrumentation and measurement services for experimenters on selected types of GENI aggregates. GIMI builds on OMF and its associated measurement library, OML [12], which are already used on GENI WiMAX testbeds. The GIMI toolset adds utilities for pushing measurements from a WiMAX testbed to a user’s personal storage on an iRODS [13] data grid, and for visualization and presentation of measurement results.

The WiMAX clients used in this work were equipped with GIMI tools, which we used to push measurements collected by an OML server on the clients to the GIMI measurement data archive (i.e. iRODS). We also used a visualization-enhanced version of the iRODS web client hosted at NYU-Poly [14] for the initial visualization and analysis of experimental results.

2) OML-enabled applications: The measurements presented in this paper were collected with existing or new applications that have been instrumented using OML libraries. During experiment runtime, the OML-enabled applications stream measurements to an OML server, which stores the measurements in a database for later retrieval or for pushing to iRODS.

To measure key experimental metrics (e.g. application-level throughput) in each experiment, we used an OML-enabled version of one of the following applications:

- The *iperf* bandwidth measurement tool [15]

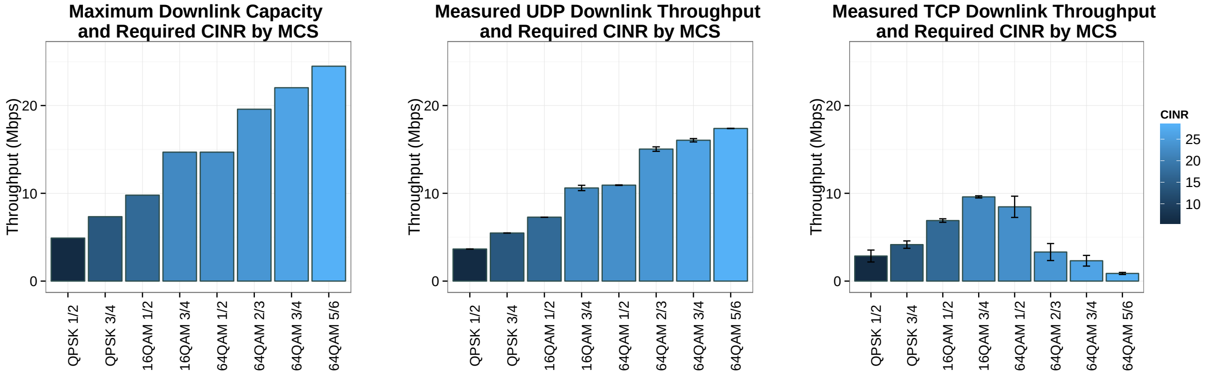

(a) (b) (c)

Fig. 2: The maximum analytical downlink capacity (2(a)), average measured UDP throughput (2(b)), and average measured TCP throughput (2(c)) for each MCS. On GENI WiMAX testbeds, when adaptive MCS is enabled, clients are assigned an MCS based on the carrier to interference plus noise ratio (CINR) of the channel. The color of each bar shows the minimum CINR that is required for the MCS to be assigned to a client. The error bars for measured UDP and TCP throughput show the standard variation across multiple trials.

- A BitTorrent client [16] based on *libtorrent-rasterbar* [17]
- A simple Python HTTP client instrumented with the *oml4py* [18] OML Python library
- An extended version of the VLC media player [16] that reports metrics from VLC's statistics module to OML.

We also collected WiMAX signal quality information from the network card and GPS location coordinates from a USB GPS device using an OML-enabled *logger* application.

III. RELATIONSHIP BETWEEN SIGNAL QUALITY AND THROUGHPUT

A key parameter of a networking testbed is the throughput available to experimental traffic. In a wireless environment, this is complicated by the relationship between wireless signal quality and throughput. Wireless systems such as WiMAX and LTE [19] use adaptive modulation and coding, where an aggressive modulation and coding scheme (MCS) which offers greater throughput can be used for a client with a strong signal, while for a client with a weak signal, a more conservative MCS must be used to maintain acceptable error rates.

Under adaptive modulation and coding, therefore, each fixed node is assigned a certain MCS based on some signal quality metric, and its download rates are effectively upper bounded by the maximum capacity available for that MCS. The upper bound on download rates for each MCS under the configuration used by the GENI WiMAX BSs in this study was calculated according to the analytical model described in [20], and the results are given in Figure 2(a).

These calculations assume a lossless channel, so actual measurements of throughput over a wireless channel are expected to be somewhat lower.

IV. EXPERIMENTS WITH FIXED WIMAX NODES

In this section, we describe measurements of downlink throughput on fixed WiMAX nodes for TCP and UDP traffic. We compare the measured throughput for TCP and UDP to the analytical maximum derived in Section III for every MCS used on GENI WiMAX testbeds.

A. Experimental Methodology

The experiments on fixed WiMAX nodes were conducted on the WITest [21] GENI WiMAX testbed at NYU-Poly. This testbed includes a set of nodes at fixed locations, physically distributed so that they see a wide range of WiMAX signal characteristics. The testbed is open to the public by reservation, and utilities from the GENI WiMAX, OMF, OML, and GIMI toolsets provide a rich environment for users.

The measurements presented in this section were collected using the OML-enabled *iperf*. Each fixed WiMAX node was configured to act as a UDP and TCP *iperf* receiver, with another host on the campus network using *iperf* to send traffic to each one in turn. At the same time, measurements of WiMAX signal characteristics, including carrier to interference plus noise ratio (CINR) for the downlink channel and the modulation and coding scheme (MCS) used on the channel, were collected from the nodes and from the WiMAX BS. The duration of each experiment was 60 seconds and every experiment was repeated five times, with the achieved data payload throughput given as the average of the five trials.

B. Results

The values of measured throughput for the UDP traffic are plotted alongside the maximum theoretical capacity in Figure 2(b). After adjusting for the estimated MAC, IP, and UDP header overhead, the UDP traffic achieves on average 76% of the upper bound on capacity. The maximum achieved throughput was 17.4 Mbps, for clients with a CINR of 28 dB or greater that can sustain 64QAM 5/6 as their MCS.

The measured throughput for the TCP traffic is given in Figure 2(c). After adjusting for the estimated MAC, IP, and TCP header overhead, TCP traffic achieved on average only 44% of the maximum analytical capacity from Section III. The maximum achieved throughput was 9.6 Mbps, for clients using 16QAM 3/4. The lowest throughput was 872 kbps for 64QAM 5/6.

The variability of throughput in time was much greater for TCP than for UDP. For TCP measurements, the standard deviation of throughput across the 60 second interval of the experiment was 42% of the mean on average, while for UDP this figure was less than 1% of the mean. This may be attributed to the cyclic behavior of TCP, which expands and contracts its congestion window in response to packet loss.

At first glance, TCP throughput appears to increase and then decrease with more aggressive MCS. However, a pattern emerges in Table II, which lists the fraction of maximum analytical capacity achieved by TCP and UDP for each MCS. We observe that for a given modulation scheme, as the code rate increases, TCP achieves less of the maximum capacity. Similarly, TCP utilization of link capacity is much less for higher order modulation (64QAM) relative to 16QAM or QPSK. This suggests that perhaps the base station is too aggressive in increasing the MCS assigned to a client.

MCS	UDP	UDP+headers	TCP	TCP+headers
QPSK 1/2	0.75	0.77	0.58	0.61
QPSK 3/4	0.75	0.77	0.57	0.59
16QAM 1/2	0.74	0.77	0.70	0.73
16QAM 3/4	0.72	0.74	0.65	0.68
64QAM 1/2	0.74	0.77	0.58	0.60
64QAM 2/3	0.77	0.79	0.17	0.18
64QAM 3/4	0.73	0.75	0.11	0.11
64QAM 5/6	0.71	0.73	0.04	0.04

TABLE II: Percent of maximum analytical capacity achieved by UDP and TCP, before and after adjusting for estimated MAC, IP, and UDP/TCP header overhead.

In every instance, UDP achieved higher throughput than TCP for the same MCS. The poor performance of a TCP flow relative to a UDP stream in saturating a link is not altogether surprising, as similar effects have been observed for cable and DSL links [22]. Furthermore, the problems affecting TCP over other wireless links are widely known [23]. TCP assumes that all losses are due to congestion and invokes congestion control mechanisms that reduce transmission rates in response to packet errors. This becomes problematic over wireless links, which suffer from high error rates. Also, the WiMAX link has a high round trip time (RTT) and also suffers from link asymmetry and occasional long delays, all factors which have been shown to be deleterious to TCP performance [24] [25].

V. EXPERIMENTS WITH MOBILE WIMAX NODES

The measurements described in this section extend the results of Section IV to consider the impact of mobility on an application using the WiMAX link. We include in this study three popular Internet applications: a traditional HTTP file download application, a BitTorrent peer-to-peer file download application, and a video streaming client. Each application is evaluated in two mobile wireless contexts: the dense urban area in which the NYU-Poly campus is located, and the rural setting of the UMass Amherst campus.

A. *Experimental Methodology*

1) Wireless Propagation Environment: NYU-Poly is located in downtown Brooklyn, within one of New York City's largest business districts. The neighborhood around the antenna consists mainly of high-rise commercial, civic, and residential buildings with a range of building heights up to 157 meters. It is a highly dynamic environment, with large moving vehicles and pedestrian traffic in the radio path. The antenna is mounted on the roof of a four-story building, at a low height relative to surrounding buildings.

The UMass Amherst campus is in a semi-rural environment, with only a few tall buildings and little traffic. The area around the base station includes flat farm land as well as some educational, residential, and commercial buildings. The antenna at UMass Amherst is situated on the roof of a 16-story building, one of the highest in the wireless coverage area.

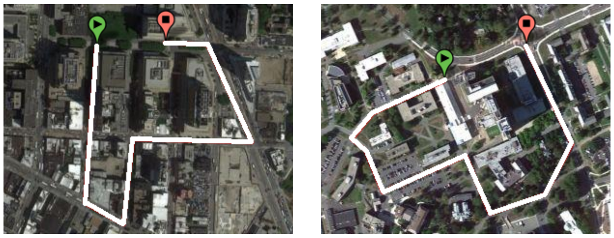
(a) (b)

Fig. 3: Mobility pattern used at NYU-Poly (3(a)) and at UMass Amherst (3(b)). Google Maps imagery ©2012 Bluesky, DigitalGlobe, GeoEye, Sanborn, USDA Farm Service Agency

2) Mobility Pattern: The measurements presented in this section were collected on a WiMAX-equipped laptop computer as it travels within the wireless coverage area. The paths are meant to be representative of a typical walking route through the area in which the BS is located.

The paths selected at NYU-Poly (see Figure 3(a)) and at UMass Amherst (see Figure 3(b)) are similar in shape, and measure 1 km in length. The WiMAX received signal strength indicator (RSSI) values along both paths, sampled at 2 m intervals, are shown in Figure 4. An average RSSI of -64 dBm was measured on both paths.

However, because of the differences in the radio propagation environment, the characteristics of the two paths, described in Table III, are quite different. The NYU-Poly path sees much greater variation in signal quality relative to the UMass Amherst path, because of the increased quantity and dynamic quality of the radio path obstructions in the dense urban environment in which it is located. More dramatic fluctuations in RSSI are observed on the urban NYU-Poly path, and it also has a higher maximum RSSI and a lower minimum RSSI.

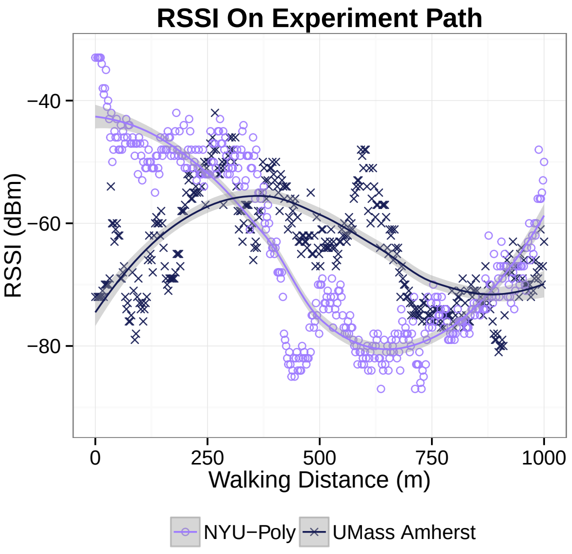

Fig. 4: Measured RSSI on experiment paths at NYU-Poly and at UMass Amherst, sampled at 2 m intervals.

	NYU-Poly	UMass Amherst
Minimum RSSI (dBm)	-87.0	-81.0
Mean RSSI (dBm)	-64.5	-64.0
Maximum RSSI (dBm)	-33.0	-42.0
SD RSSI (dBm)	14.3	8.6

TABLE III: Mobile Signal Quality Characteristics

3) Measurement Applications: Three applications are considered in this study of mobile application performance. The applications were selected because they represent a diverse set of requirements from the network.

First, we evaluate an ordinary HTTP download using a simple Python-based HTTP client to download a large DVD image of a popular Linux distribution from a server located on the campus network. The client logs the number of bytes downloaded and the timestamp each time the client downloads a chunk of 8192 bytes, from which we calculate the download rate (smoothed over a 1 second interval).

Next, we use a BitTorrent client to download the same DVD image from a torrent in which hundreds of peers on the public Internet are participating. The BitTorrent client is instrumented to record the current download rate at regular intervals.

Finally, we consider a UDP video streaming scenario, in which we use the open-source VLC media player to stream video from a host on the campus network to a mobile WiMAX client. The video was encoded with parameters that are typical for streaming video to a mobile device over the Internet (MPEG-4 video encoded at 700kpbs with a 24fps frame rate and 854x480 resolution). At the client, the video was buffered for 300ms to smooth jitter. Quality metrics including video bitrate were recorded at regular intervals.

B. Results

The results of these experiments show that the impact of mobility on application performance depends on the characteristics of the application and of the wireless setting.

Based on the TCP results from Section IV, we expect that an HTTP download - which is essentially a single TCP session - will not see good data rates over the WiMAX link. This is confirmed by the measurement results shown in Figure 5. The HTTP file download does not saturate the WiMAX link. During periods of especially low RSSI, the throughput is temporarily reduced to zero, and it takes some time to recover to a steady flow. The highest download rates of up to 8 Mbps are seen during periods where the RSSI is constant at a moderate value, at which time the client is likely to be assigned either 16QAM 3/4 or 64QAM 1/2 as its download MCS. This is consistent with the results described in Section IV, where the highest measured throughput for TCP in the fixed setting was close to 9 Mbps for 16QAM 3/4 and 64QAM 1/2.

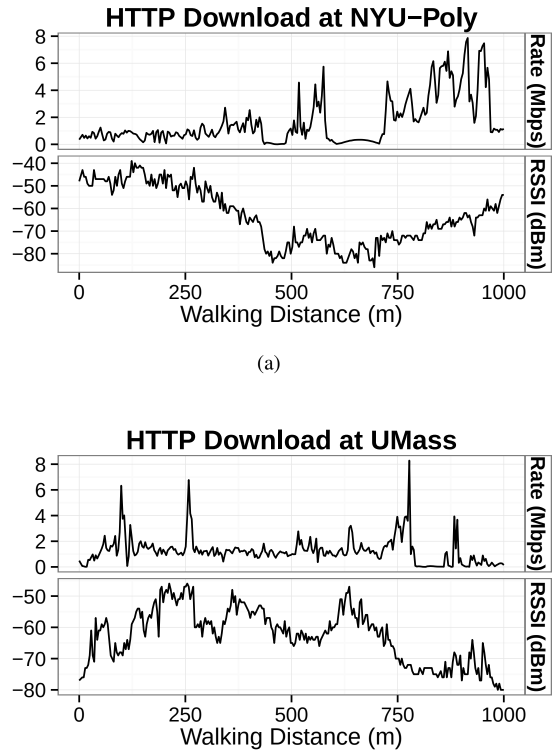

(a)

(b)

Fig. 5: HTTP download rate and RSSI at NYU-Poly (5(a)) and at UMass Amherst (5(b)).

For the HTTP download, the average download rate was slightly greater at NYU-Poly than at UMass Amherst (1.63 Mbps and 1.24 Mbps, respectively). The standard variation of the download rate was much greater at NYU-Poly (1.76 Mbps vs. 1.05 Mbps at UMass Amherst), which is consistent with the higher variability in RSSI seen in the urban setting.

In contrast to an ordinary HTTP file transfer, which uses a single TCP session between a client and a server, a BitTorrent transfer establishes multiple TCP sessions between a client and

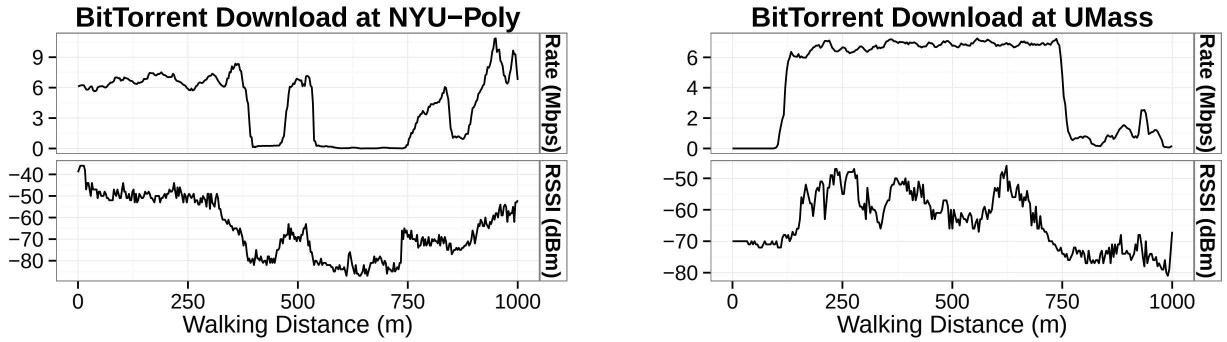

(a)

(b)

Fig. 6: BitTorrent download rate and RSSI at NYU-Poly (6(a)) and at UMass Amherst (6(b)).

a number of peers. There exists a great body of work showing that striping data transfers across parallel TCP connections, as BitTorrent does, can yield higher throughput [26] [22], especially in lossy or high-latency, high-bandwidth environments. Utilizing multiple TCP sessions reduces the negative effects of random packet loss, since losses are distributed among the TCP flows and at any given time, the likelihood that a flow will invoke congestion avoidance due to packet loss is lower. This increases aggregate throughput and allows the application to utilize the wireless link more efficiently.

The BitTorrent measurement results are shown in Figure 6. We can see by inspection that the shape of the download rate curve closely mimics the shape of the RSSI curve, with higher download rates corresponding to higher RSSI values. This is consistent with our expectation that the BitTorrent download will be more efficient at saturating the wireless link. The shape of the BitTorrent curves are smooth relative to the shape of the HTTP download curves in Figure 5, i.e. the BitTorrent application is less sensitive to changes in RSSI.

Overall, the average BitTorrent download rate (4.25 Mbps at NYU-Poly and 4.51 Mbps at UMass Amherst) is more than double the equivalent HTTP download rate. Although the variation in download rate is higher at NYU-Poly than at UMass Amherst (3.12 Mbps vs. 2.96 Mbps), the difference is not as dramatic as in the HTTP case, indicating that the effects of the dynamic RSSI in the urban setting are not as severe.

In the video streaming experiment, the video bitrate (700 kbps) is well below the maximum capacity of the link, and is not likely to saturate the link except under the poorest signal conditions. The video bitrate seen along the mobile path is shown in Figure 7. The mobile clients at both NYU-Poly and UMass Amherst saw bit rates of approximately 700 kbps. However, when the client moves into areas with a weak signal, the received video bitrate becomes more irregular, resulting in blocky, choppy video. This effect may be seen, for example, in the interval between 450 m and 750 m in Figure 7(a).

Because the video bitrate was well below link capacity, video quality remains acceptable throughout. However, we still see a greater variation in quality at NYU-Poly (89 kbps) relative to UMass Amherst (78 kbps). This effect - and a corresponding decline in user experience - would most likely become more pronounced with increasing video bitrate.

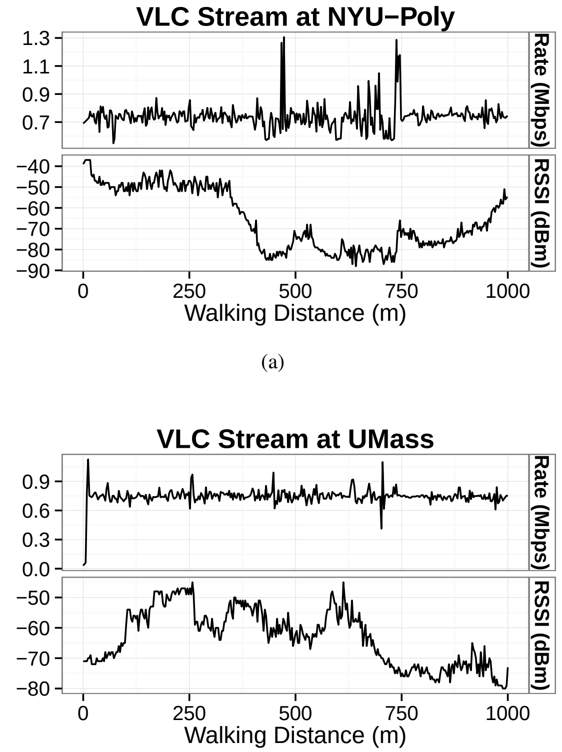

(a)

(b)

Fig. 7: VLC video data rate and RSSI at NYU-Poly (7(a)) and at UMass Amherst (7(b)).

The measurements from the three mobile experiments are plotted alongside one another in Figure 8. From this figure, we can observe some key characteristics of the WiMAX network performance. We see that BitTorrent throughput was much greater than HTTP throughput at both sites. Although average throughput was similar for the HTTP experiment and the video experiment, the UDP video traffic experienced a much

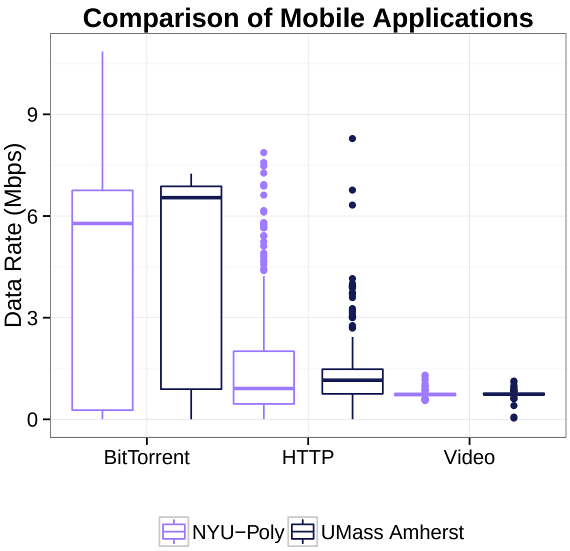

Fig. 8: Measurements from mobile experiments. The boxplot shows the mean data rate for each experiment, with the upper and lower hinges corresponding to the first and third quartiles. Outliers beyond 1.5 interquartile range (IQR) of the hinges are plotted as points.

more consistent data rate than the TCP HTTP experiment. The spread of data rates was wider at NYU-Poly for every experiment, which is consistent with the wider spread of RSSI values in the urban environment. Also, when means are calculated without the influence of outliers, NYU-Poly sees worse average performance for every experiment, suggesting that the urban environment results in lower average throughput in addition to greater variability.

VI. RELATED WORK

To date, we are the first to systematically evaluate a wide range of applications on identical mobile WiMAX networks in dramatically different wireless settings.

However, there has been extensive work towards WiMAX performance measurements. Some of this work describes simulation results, or measurements on networks using the older "fixed WiMAX" standard. For example, [27] and [28] analyze the performance of video over mobile WiMAX networks using simulations. [29] [30] [31] describe field measurements from experiments carried out over research testbeds.

More recent studies used mobile WiMAX networks and assessed the performance of wireless Internet applications including BitTorrent [32], World of Warcraft [33], TCP and UDP [34], TCP [35], and VoIP [36] [37].

However, these experiments were conducted over commercial WiMAX networks and as such have been subject to the influence of carrier scheduling policies, per-user bandwidth caps, and competing traffic flows. For example, the WiMAX service provider used in [35] limits the maximum data rates to 1.5 Mbps downlink and 256 Kbps uplink, and midway through the experiment the authors saw a dramatic change in measured throughput which they attribute to changes in the service provider's traffic classification policies. Because the commercial network must be treated as a "black box", it is unclear to what extent the findings of these experiments will apply to another network. We therefore cannot directly compare measurements from different studies in any meaningful way. In our experiments, we have full control of the WiMAX network, including radio control parameters at the BS, and thus we can ensure that they are consistent throughout the experiment and between experiment sites. Furthermore, we can isolate the effect of network-specific parameters, such as adaptive modulation and coding assignments, which has not been possible for the work conducted on commercial networks.

Other measurements studies have investigated the performance of different kinds of cellular data networks. Notably, [38] describes the throughput and other performance characteristics of commercial LTE networks using measurements from end-user devices.

VII. CONCLUSION

The goal of this work was to gain a greater understanding of the behavior of Internet applications on the GENI WiMAX platform. This work has the potential to aid other researchers in designing realistic experiments and interpreting their results in the context of the experimental setting.

In this paper, we described a thorough investigation of the behavior of popular Internet applications on a GENI WiMAX testbed in different wireless settings. Through the systematic collection and analysis of empirical measurements, we have identified characteristics of a traffic flow or a wireless network setting that affect application performance in a meaningful way. For example, we have shown that:

- The maximum throughput achieved by an application depends on the quality of the wireless signal it receives. However, there is not a direct relationship between wireless signal quality and throughput, and a better wireless channel does not necessarily yield improved throughput.
- A single TCP flow is not resilient to packet loss, and will not achieve the full throughput that is possible given the capacity of the link. A UDP stream or a set of parallel TCP flows will see better throughput.
- The broader context of the wireless setting matters in mobile experiments. In an urban area, where signal quality is more dynamic, the performance of the wireless network will be highly dynamic as well. For mobile experiments in an urban area, applications that require a consistent download rate (such as multimedia streaming applications) may see poor performance.

As future work, we intend to collect additional measurements to verify the conclusions suggested by the preliminary measurements described in this paper, e.g. more measurements for each path and new measurements on additional paths. We would also like to collect measurements from other kinds of Internet applications, such as a VoIP application and a gaming application. For each application under consideration, we will consider additional quality metrics in addition to throughput, such as the number of peers for BitTorrent transfers and the number of corrupted frames for the video streaming application. We may then perform a deeper analysis of the results, and gain a better understanding of what classes of applications will see "good" performance on a WiMAX network and what classes are likely to perform poorly.

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant No. CNS-0714770 and under Graduate Research Fellowship Program Grant No. 1104522.

REFERENCES

[1] D. Kotz, C. Newport, R. S. Gray, J. Liu, Y. Yuan, and C. Elliott, "Experimental evaluation of wireless simulation assumptions," in *Proceedings of the 7th ACM international symposium on Modeling, analysis and simulation of wireless and mobile systems*, MSWiM '04, (New York, NY, USA), p. 7882, ACM, 2004.

[2] K. Pawlikowski, H.-D. Jeong, and J.-S. Lee, "On credibility of simulation studies of telecommunication networks," *IEEE Communications Magazine*, vol. 40, pp. 132 –139, Jan. 2002.

[3] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the orbit radio grid testbed for evaluation of next-generation wireless network protocols," in *Wireless Communications and Networking Conference, 2005 IEEE*, vol. 3, pp. 1664 – 1669 Vol. 3, March 2005.

[4] M. Muhammad and J. Cappos, "Towards a representive testbed: Harnessing volunteers for networks research," in *The First GENI Research and Educational Workshop, GREE*, vol. 12, 2012.

[5] *Global Environment for Network Innovation*. http://www.geni.net/.

[6] "IEEE standard for Local and Metropolitan Area Networks part 16: Air interface for fixed and mobile broadband wireless access systems amendment 2: Physical and medium access control layers for combined fixed and mobile operation in licensed bands and corrigendum 1," *IEEE Std 802.16e-2005 and IEEE Std 802.16-2004, Cor 1-2005 (Amendment and Corrigendum to IEEE Std 802.16-2004)*, pp. 0–1–822, 2006.

[7] *Large-scale GENI Instrumentation and Measurement Infrastructure*. http://gimi.ecs.umass.edu/.

[8] T. Rakotoarivelo, M. Ott, G. Jourjon, and I. Seskar, "OMF: a control and management framework for networking testbeds," *SIGOPS Oper. Syst. Rev.*, vol. 43, pp. 54–59, Jan. 2010.

[9] G. Bhanage, I. Seskar, R. Mahindra, and D. Raychaudhuri, "Virtual basestation: architecture for an open shared wimax framework," in *Proceedings of the second ACM SIGCOMM workshop on Virtualized infrastructure systems and architectures*, VISA '10, (New York, NY, USA), pp. 1–8, ACM, 2010.

[10] S. Hanks, T. Li, D. Farinacci, and P. Traina, "Generic Routing Encapsulation (GRE)." RFC 1701 (Informational), Oct. 1994.

[11] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The Click modular router," *ACM Trans. Comput. Syst.*, vol. 18, pp. 263–297, Aug. 2000.

[12] O. Mehani, G. Jourjon, T. Rakotoarivelo, and M. Ott, "An instrumentation framework for the critical task of measurement collection in the future internet," tech. rep., Nicta, Eveleigh, Sydney, NSW, Australia, October 2012.

[13] *IRODS*. https://www.irods.org/.

[14] *Extended IRODS web client*. https://www.witest.poly.edu/irods/web/.

[15] *OML-Instrumented Applications (NICTA)*. http://omlapp.mytestbed.net/projects/omlapp/wiki/OML-instrumented_Applications.

[16] *OML Enabled Applications (NYU-Poly)*. http://witestlab.poly.edu/index.php/wirelesstestbeds/omfnettest/oml-applications.html.

[17] *libtorrent rasterbar*. http://www.rasterbar.com/products/libtorrent/.

[18] *oml4py: OML client for Python*. http://pypi.python.org/pypi/oml4py/.

[19] *3GPP-LTE*. http://www.3gpp.org/LTE/.

[20] C. So-In, R. Jain, and A.-K. Tamimi, "Capacity evaluation for IEEE 802.16e mobile WiMAX," *J. Comp. Sys., Netw., and Comm.*, vol. 2010, p. 1:11:12, Jan. 2010.

[21] *Wireless Implementation Testbed Laboratory*. http://witestlab.poly.edu/.

[22] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescap, "Broadband internet performance: a view from the gateway," *SIGCOMM-Computer Communication Review*, vol. 41, no. 4, p. 134, 2011.

[23] G. Xylomenos, G. Polyzos, P. Mahonen, and M. Saaranen, "Tcp performance issues over wireless links," *Communications Magazine, IEEE*, vol. 39, pp. 52 –58, Apr. 2001.

[24] H. Balakrishnan, V. Padmanabhan, and R. Katz, "The effects of asymmetry on tcp performance," *Mobile Networks and Applications*, vol. 4, pp. 219–241, 1999.

[25] A. Gurtov, "Effect of delays on tcp performance," in *Emerging Personal Wireless Communications* (O. Martikainen, J. Porras, and J. Hyvrinen, eds.), vol. 67 of *IFIP International Federation for Information Processing*, pp. 87–105, Springer US, 2002.

[26] T. Hacker, B. Athey, and B. Noble, "The end-to-end performance effects of parallel tcp sockets on a lossy wide-area network," in *Parallel and Distributed Processing Symposium., Proceedings International, IPDPS 2002, Abstracts and CD-ROM*, pp. 434–443, IEEE, 2002.

[27] H. Chughtai, S. Malik, and M. Yousaf, "Performance evaluation of transport layer protocols for video traffic over WiMAX," in *Multitopic Conference, 2009. INMIC 2009. IEEE 13th International*, pp. 1 –6, Dec. 2009.

[28] H.-H. Juan, H.-C. Huang, C. Huang, and T. Chiang, "Scalable video streaming over mobile WiMAX," in *Circuits and Systems, 2007. ISCAS 2007. IEEE International Symposium on*, pp. 3463 –3466, May 2007.

[29] K. Pentikousis, J. Pinola, E. Piri, and F. Fitzek, "An experimental investigation of VoIP and video streaming over fixed WiMAX," in *Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks and Workshops, 2008. WiOPT 2008. 6th International Symposium on*, pp. 8 –15, April 2008.

[30] K. Filis, E. Theodoropoulou, and G. Lyberopoulos, "The effect of a rapidly changing urban environment on nomadic WiMAX performance," in *Mobile and Wireless Communications Summit, 2007. 16th IST*, pp. 1 –5, July 2007.

[31] F. Yousaf, K. Daniel, and C. Wietfeld, "Performance evaluation of IEEE 802.16 WiMAX link with respect to higher layer protocols," in *Wireless Communication Systems, 2007. ISWCS 2007. 4th International Symposium on*, pp. 180 –184, Oct. 2007.

[32] X. Wang, S. Kim, T. T. Kwon, H.-c. Kim, and Y. Choi, "Unveiling the BitTorrent performance in mobile WiMAX networks," in *Proceedings of the 12th international conference on passive and active measurement*, PAM'11, (Berlin, Heidelberg), pp. 184–193, Springer-Verlag, 2011.

[33] X. Wang, H. Kim, A. Vasilakos, T. Kwon, Y. Choi, S. Choi, and H. Jang, "Measurement and analysis of World of Warcraft in mobile WiMAX networks," in *Network and Systems Support for Games (NetGames), 2009 8th Annual Workshop on*, pp. 1 –6, Nov. 2009.

[34] D. Kim, H. Cai, M. Na, and S. Choi, "Performance measurement over mobile WiMAX/IEEE 802.16e network," in *World of Wireless, Mobile and Multimedia Networks, 2008. WoWMoM 2008. 2008 International Symposium on a*, pp. 1 –8, June 2008.

[35] E. Halepovic, Q. Wu, C. Williamson, and M. Ghaderi, "TCP over WiMAX: A measurement study," in *Modeling, Analysis and Simulation of Computers and Telecommunication Systems, 2008. MASCOTS 2008. IEEE International Symposium on*, pp. 1 –10, Sept. 2008.

[36] M. Han, Y. Lee, S. Moon, K. Jang, and D. Lee, "Evaluation of VoIP quality over WiBro," in *Proceedings of the 9th international conference on passive and active network measurement*, PAM'08, (Berlin, Heidelberg), pp. 51–60, Springer-Verlag, 2008.

[37] E. Halepovic, M. Ghaderi, and C. Williamson, "Multimedia application performance on a WiMAX network," pp. 725309–725309, Jan. 2009.

[38] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck, "A close examination of performance and power characteristics of 4G LTE networks," in *Proceedings of the 10th international conference on Mobile systems, applications, and services*, MobiSys '12, (New York, NY, USA), p. 225238, ACM, 2012.