Performance Characteristics of an Operational WiMAX Network

James Martin, *Member, IEEE,* and James Westall, *Affiliate, IEEE CS*

Abstract—The term WiMAX is used to refer to a collection of standards, products, and service offerings derived from the IEEE 802.16 family of standards for metropolitan area wireless networks. A significant body of published research in the WiMAX domain exists, but the focus of much of it is on the use of analytic or simulation models to evaluate aspects of physical layer protocols, medium access control protocols, or proposed scheduling algorithms. It this paper we describe performance characteristics of an operational WiMAX testbed upon which we were able to conduct controlled experiments in the absence of competing traffic. We characterize latency, throughput, protocol overhead, packet loss, and the impact of WiMAX on TCP dynamics.

Index Terms—WiMAX, IEEE 802.16, wireless, performance.

1 Introduction

The term WiMAX, an acronym for Worldwide Interoperability for Microwave Access, is commonly used to refer to a collection of standards, products, and service offerings derived from the IEEE 802.16 family of standards [1], [2]. The IEEE standards include many implementation options that are left to equipment vendors. Conflicting design choices can make interoperation of the equipment of multiple vendors problematic. The WiMAX forum was organized by equipment vendors in 2001 to define operational profiles, certify interoperability, and promote the use of the technology. A discussion of the roles of the IEEE and the WiMAX forum in the development of the standards and profiles can be found in [3].

The equipment described in this paper is compliant with the IEEE 802.16-2004 standard [1] which is sometimes called IEEE 802.16d or "fixed WiMAX" because it does not support seamless handoff for mobile clients. The subsequent amendment, IEEE 802.16e-2005 [2], sometimes called 802.16e, added support for mobile clients. Perspectives on the evolution of WiMAX may be found in [4], [5], [6]. A very thorough discussion of the WiMAX physical layer is provided in [4]. A discussion of WiMAX as it relates to alternative wireless technologies is found in [7].

At present WiMAX usage is not widespread when compared to competing access network technologies. In urban areas it competes with WiFi mesh technology, and in both urban and suburban areas it competes with cable and DSL service. Although originally designed to provide metropolitan area network (MAN) service, its main competitive advantage against competing technologies lies in sparsely populated rural areas where the cost of providing high bandwidth fiber coverage is prohibitive.

• *The authors are with the School of Computing, Clemson University, Clemson, SC 29634.*
E-mail: see http://www.cs.clemson.edu

Where commercial offerings do exist, the service providers are generally unwilling to release usage and performance data. Furthermore, it is very difficult to conduct controlled experiments that measure best case throughput and latency on an operational public network. Therefore, virtually all published studies of the performance characteristics of WiMAX systems have been derived from simulation or analytic models.

The objective of this paper is to augment the results obtained in simulation studies of hypothetical equipment with measured results obtained from an operational WiMAX testbed. Our focus is upon characterizing the latency, throughput, and overhead properties of the network. The important issue of coverage will be addressed in a subsequent paper.

The testbed is deployed on the campus of Clemson University and uses M/A-COM's Vida WiMAX equipment. The network operates in the 4.9 gigahertz (GHz) public safety band which is comprised of ten channels of five MHz each spanning 50 MHz of spectrum between 4940 and 4990 MHz. This spectrum was allocated by the Federal Communications Commission (FCC) in 2002 for fixed and mobile wireless services in support of public safety [8]. Both base and subscriber stations operating in this spectrum are limited to no more than 27 dBm of transmitter power and no more than 40 dBm of effective isotropic radiated power. Although a WiMAX Forum profile for 4.9 GHz has not yet been defined, WiMAX equipment vendors have agreed on a set of operating parameters allowing interoperability. These parameters are consistent with the 802.16-2004 standard and are used in equipment currently offered by Airspan, M/A-COM, and Nortel.

The authors of this paper were funded by the National Institute of Justice to evaluate the suitability WiMAX equipment operating in this spectrum for use in public safety applications. The remainder of the paper reports on aspects of that study. The results reported in this paper were obtained during an evaluation period in

which the network was under the operational control of the authors and operating under a license held by the Clemson University Police Department and the City of Clemson Police Department.

2 The WiMAX testbed

The equipment used in the study includes a M/A-COM VIDA Broadband MAVM-VMXBD hardened base station, M/A-COM VIDA Broadband MAVM-VMCLL subscriber stations, and Airspan EasyST subscriber stations. In the remainder of this section we review aspects of the 802.16-2004 standard that pertain to this equipment and this study.

2.1 The Physical Layer

The 802.16-2004 standard defines single carrier (SC), orthogonal frequency division multiplexing (OFDM), and orthogonal frequency division multiple access (OFDMA) modes of operation at the physical layer. The M/A-COM equipment implements only OFDM operation.

Operational parameters that bound the capacity of an OFDM WiMAX physical layer include:

- channel bandwidth,
- number of data-carrying subchannels,
- modulation and forward error correction (FEC) technique and,
- duplexing mode (time or frequency division duplexing).

The M/A-COM equipment can operate in any one of the ten channels of the 4.9 GHz spectrum. The 5 Mhz bandwidth is partitioned into 256 subchannels as specified in the standard: eight pilot channels are used in physical layer synchronization; 55 channels are used as guard bands; and 192 channels carry data. A null carrier is transmitted on the remaining center frequency channel. The OFDM timing data used by the M/A-COM equipment is summarized in the terminology of p. 428 of the standard in Table 1.

TABLE 1
OFDM Timing Data

Parameter	Value
Bandwidth	$5.000 \times 10^{+6}$
N_{fft}	$2.560 \times 10^{+2}$
Sampling frequency	$5.760 \times 10^{+6}$
Subcarrier spacing	$2.250 \times 10^{+4}$
Useful symbol time	4.444×10^{-5}
Cyclic prefix time	5.556×10^{-6}
OFDM Symbol time	5.000×10^{-5}
Frame time	1.000×10^{-2}
Symbols / frame	$2.000 \times 10^{+2}$

The sampling frequency is computed as $\lfloor n \times Bandwidth/8000 \rfloor * 8000$ where $n = 144/125$ for a channel whose bandwidth is an even multiple of 1.25 Mhz. The subcarrier spacing is the sampling frequency divided by the total number of subchannels (FFTs). The useful symbol time is the inverse of the subcarrier spacing. The cyclic prefix time is the useful symbol time divided by 8, and the OFDM symbol time is the sum of the useful symbol time and the cyclic prefix time.

The data carrying capacity of each symbol is a function of the number of data carrying subchannels, the modulation technique, and the number of bits reserved for forward error correction. Values supported by 802.16-2004 are shown in Table 2. The column labeled bits per sample shows the aggregate number of data and FEC bits that can be encoded on a single channel using the given modulation scheme during a single OFDM symbol time. The number of data bits per symbol is obtained by multiplying bits per sample by the number of data channels (192) and the coding rate (fraction of bits representing data) shown in the leftmost column. The number of symbols per network layer protocol data unit (NPDU) is 1500 divided by the number of bytes per symbol.

TABLE 2
Symbol Capacity

Modulation	Bits per Sample	Data Bits per Sym	Kbps	Syms per NPDU
BPSK 1/2	1	96	960	125
QPSK 1/2	2	192	1920	62.5
QPSK 3/4	2	288	2822	41.67
16-QAM 1/2	4	384	3840	31.25
16-QAM 3/4	4	576	5760	20.83
64-QAM 2/3	6	768	7680	15.63
64-QAM 3/4	6	864	8640	13.89

The modulation technique may change dynamically in response signal quality. In the M/A-COM equipment, modulation changes are triggered by changes in the carrier to interference and noise ratio (CINR). CINR levels needed to trigger change are configurable by the system administrator. A dual level triggering mechanism is used to prevent "modulation flapping."

The standard defines an optional MAC level ARQ mechanism designed to provide fast recovery from physical layer errors not corrected by the FEC. ARQ is not implemented in our M/A-COM equipment.

The M/A-COM equipment employs time division duplexing (TDD). A single transmit frequency is used, and the equipment rapidly alternates between transmit and receive modes. Frame time is 10 ms which yields 200 physical layer symbols per frame. A frame is comprised of a downlink subframe in which the base station transmits and the subscribers receive followed by an uplink subframe in which the reverse is true. Each subframe may be further subdivided into transmission *bursts* with modulation and/or coding rate changing dynamically from burst to burst within a single subframe. Relative lengths of uplink and downlink subframes are configurable. We employed a nominal 50/50 split but later discovered that downlink MAC overhead was substantially larger than on the uplink, and the resulting split at the network layer was actually closer to 44/56.

The upper bound on physical layer throughput in

Kbps is shown in Table 2. It assumes a 50/50 downlink-uplink split. Throughput values are obtained by multiplying bytes per symbol by 100 symbols per subframe by 100 frames per second. Because of PHY and MAC overhead, actual network layer throughput is considerably smaller. Sources of symbol consuming overhead at both the PHY and MAC layers are identified on p. 449 of the standard. Transmit-receive and receive-transmit transition gaps (TTG, RTG) are required between subframes. The downlink subframe begins with a two-symbol long preamble, and the uplink subframe employs a one-symbol short preamble. The downlink MAC layer data begins with the frame control header (FCH) which is is always transmitted using BPSK 1/2. It contains the downlink frame prefix (DLFP) which contains the burst profiles of up to the first three bursts of the downlink subframe. A burst profile contains the starting offset measured in symbols[1] and the modulation technique used in the burst.

The first burst following the FCH is sometimes called the broadcast burst. It is transmitted using the least robust modulation technique that all subscribers are thought to be able to presently decode. This burst always contains the uplink MAP (UL-MAP) which describes the allocation of of uplink symbols. In the M/A-COM implementation the UL-MAP in frame n describes the uplink symbol allocation in frame $n + 1$. In the default configuration the M/A-COM implementation always includes downlink and uplink channel descriptors (DCD, UCD) and a downlink map (DL-MAP) in the broadcast burst, but the frequency at which the DCD and UCD are sent is configurable. If a downlink subframe contains more than three bursts, the DL-MAP carries their burst profiles.

The precise number of overhead symbols per frame is not directly specified in the standard because it depends upon both configuration (e.g., how often the DCD and UCD are sent) and provisioning (e.g., the number of active RTPS flows). Nevertheless, if the maximum network layer throughput is known, then the average number of overhead symbols per subframe may be inferred as described later in this paper.

2.2 WiMAX provisioning

2.2.1 Service flows

Unlike WiFi networks, WiMAX networks support very fine grained control over provisioning of network traffic flows. An individual flow is a unidirectional entity referred to as a $service flow$. A three-phase activation model for service flows is described on p. 223 of the standard.

In the $provisioned$ phase, the service flow is known to the base and subscriber stations by a 32-bit service flow identifier (SFID), but no network resources are allocated to it. $Admission$ and $activation$ may be triggered by either the base station or the subscriber using a protocol called dynamic service activation (DSA). Admission may fail if insufficient resources to meet QoS requirements are unavailable. If admission succeeds, the service flow is activated and network capacity is allocated to it. Traffic associated with an active service flow is known by a 16 bit connection identifier (CID). This aspect of the standard is designed to support transient connections such as VoIP calls.

The protocols through which the subscriber station and end systems agree to trigger a dynamic service activation are not specified, and, while dynamic activation by the base station is a required element of the standard, dynamic activation capability in the subscriber is optional. Therefore, in the minimalist, but compliant, implementation of the M/A-COM equipment, provisioning, admission, and activation are simultaneous and occur at boot time. Dynamic changes are triggered only by actions of the system administrator. In this situation, provisioned but inactive service flows may wastefully consume network capacity while carrying no traffic at all.

2.2.2 Scheduling

Traffic scheduling in a WiMAX network is similar to scheduling in a DOCSIS cable network. Allocation of transmission opportunities for both downlink and uplink traffic flows is vested in the base station.

The DL-MAP and UL-MAP data structures, which contain starting offset and encoding of each burst, provide the mechanism by which the results of the scheduling policies are made known, but the the details of scheduling policies themselves are not defined in the standard and are left to the equipment vendor. M/A-COM's scheduling algorithms are proprietary and were not revealed to us.

The standard does distinguish four distinct scheduling types that pertain to the allocation of $uplink$ capacity. For Unsolicited Grant Service (UGS), periodic grants of sufficient capacity to carry the provisioned bit rate are conveyed in the UL-MAP to each subscriber station provisioned with a UGS flow. In Real-time Polling Service (RTPS) and non-Real-time Polling Service (nRTPS) flows, the UL-MAP periodically identifies dedicated slots in which the subscriber station can make contention free requests for uplink capacity. Contention slots, also identified in the UL-MAP, may be used by subscriber stations to request capacity for both Best Effort (BE) and nRTPS flows. Contention requests can be destroyed by collisions among competing subscriber stations. Collisions trigger a binary exponential backoff. As in DOCSIS, contention is minimized by permitting a subscriber station with backlogged best effort traffic to "piggyback" a request for additional capacity onto the packet currently being transmitted.

QoS attributes that may be assigned to service flows are identified on p. 695 of the standard. They include

1. Symbol based offsets are not used in either SC or OFDMA

traffic priority, maximum sustained traffic rate, maximum traffic burst, minimum reserved traffic rate, minimum tolerable traffic rate, maximum latency (delay within subscriber or base station from the time a packet is received on the wire interface until it is transmitted on the RF interface), and tolerated jitter. The standard does not specify how the QoS attributes are to be incorporated into the underlying scheduling algorithms nor does it specify any required behavior in the event of overcommitment of resources.

Two additional MAC layer capabilities also found in DOCSIS facilitate scheduling and reduce overhead. *Fragmentation* allows a large NPDU to be broken up and carried in multiple MAC layer PDUs. Fragmentation can be used to ensure that available capacity in a frame is not wasted because it is not sufficient to hold a full NPDU. *Concatenation* allows a single MAC layer PDU to carry multiple small NPDUs thus reducing MAC layer overhead. The M/A-COM equipment supports both fragmentation and concatenation.

2.3 Provisioning M/A-COM equipment

The M/A-COM base station employs an an Intel IXP2350 network processor as its main processor. It runs an ADEOS (Adaptive Domain Environment for Operating Systems) real-time variant of the 2.4.20 Linux kernel. The base station supports remote login via *telnet* and contains an NFS client that is useful in logging to persistent storage performance data generated by base station utility programs. Although the system administrator has root level access on the base station, the utilities themselves are undocumented, and mechanisms for viewing or making direct changes to provisioning while logged into the base station are not provided.

Instead, provisioning the M/A-COM equipment is accomplished using M/A-COM's web-based Unified Administration System (UAS) which runs on an auxiliary Linux computer. The provisioning data is distributed to the WiMAX base station via SNMP. The base station then conveys provisioning information to the subscribers over the air link.

The Airspan and M/A-COM subscriber stations have a closed architecture. They support a Web interface that provides the current received signal strength and uplink and downlink modulation modes, but it supports neither viewing nor modifications of provisioning data.

Provisioning a network with UAS requires defining and configuring a hierarchy of four entity types: base stations; subscriber stations; service flows; and classifier rules. An instance of each entity type has a name and some attributes. Elements higher in the hierarchy bind to entities in lower levels by using the name of the lower level entity.

2.3.1 Classifier rules

The classifier rules lie at the bottom of the hierarchy. The primary attribute of each rule is the value of a header element in an application data packet that can be used to associate the data packet with the classifier rule. Header elements that are currently supported include: the source or destination MAC address, IP address, and IP subnet address. At present classifier rules must be unidirectional and specify only a source or destination address. The classifier rules have associated priorities that govern the order in which they are applied to each packet.

2.3.2 Service flows

QoS attributes of a network flow are specified in a service flow definition. Service flows are by definition unidirectional, and so any bidirectional network flow requires two of them. Attributes of service flow definitions include the names of associated classifier rules. During network operation, an individual packet is mapped to a service flow when a header element in the packet matches a classifier rule associated with the service flow.

It is logically the case that multiple traffic classifier rules could be bound to a single service flow definition, and the UAS tool supports this. Nevertheless, M/A-COM recommended a one-to-one binding of classifier rules and service flow definitions. Other attributes of a service flow include its direction (uplink/downlink) and scheduling type (UGS, RTPS, nRTPS, or BE). Scheduling type must be specified for downlink flows, but its effect is not defined and not apparent. Most of the QoS attributes defined in the standard may be specified. The maximum bit rate attribute has been verified to impose a reliable upper bound on throughput in either direction.

2.3.3 Subscriber stations

The attributes of a subscriber station definition are its service flows. Each service flow must be assigned to at most one subscriber station at a time, and upstream and downstream flows must be properly paired together.

The Airspan and M/A-COM subscriber stations do not support dynamic service activation. Therefore, provisioned UGS, RTPS, and nRTPS service flows consume network capacity even when no application traffic is present. Each data grant or polling opportunity consumes a minimal amount of capacity in the DL-MAP. Each polling opportunity consumes two symbols in the uplink subframe, but a UGS grant consumes the full amount of provisioned uplink capacity. The standard does permit a subscriber station to assign uplink capacity to other flows when not used by the flow to which it was granted to other flows, and the M/A-COM system does this.

2.3.4 Base stations

The primary attributes of a base station are the names of its subscriber stations and the name of its base station profile. The base station profile allows one to specify physical layer parameters such as CINR thresholds that trigger changes in modulation and received signal strength indication (RSSI) thresholds that trigger power level changes.

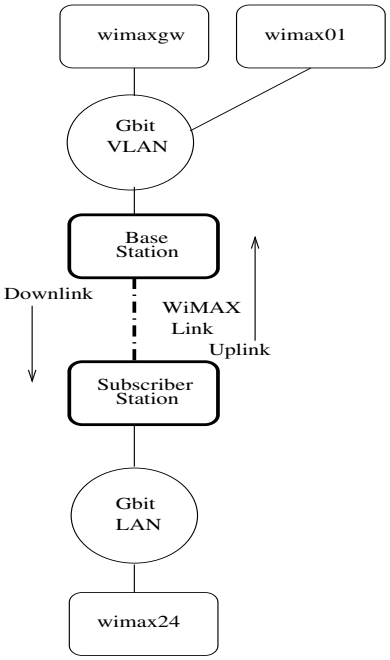

Fig. 1. WiMAX network configuration

The physical layer parameters are far less intuitive than those dealing with service flows, and errors in defining physical layer parameters can render an entire network inoperative. Fortunately, UAS provides a conservative "Default" profile that can be cloned and safely used for experimentation

3 Latency

In this section we characterize the latency experienced by unsolicited grant service (UGS), real time polling service (RTPS), and best effort (BE) service classes when *ping* type probe packets are sent both uplink and downlink across the WiMAX network. We show how the measurements obtained can be used to understand aspects of the underlying packet scheduling algorithms. Elements of the WiMAX network used in the study are shown in Fig. 1. The systems named *wimaxgw* and *wimax24* are both multi-homed Linux hosts in the School of Computing at Clemson University. In addition to the WiMAX network, these two hosts are also connected to the School of Computing's gigabit LAN.

While the tests were being conducted, three subscriber stations were provisioned and powered on. One of these, which was not involved in the test, was provisioned with a single RTPS flow. Thus the UL-MAP in each downlink subframe described a reservation slot in the next uplink subframe in which a bandwidth request could be made. However, during the measurement period no competing applications generated traffic.

Measurements were collected using a UDP client and server pair. The client periodically sends a small probe packet to the server located on the other side of the WiMAX network. The probe packet is carried in an

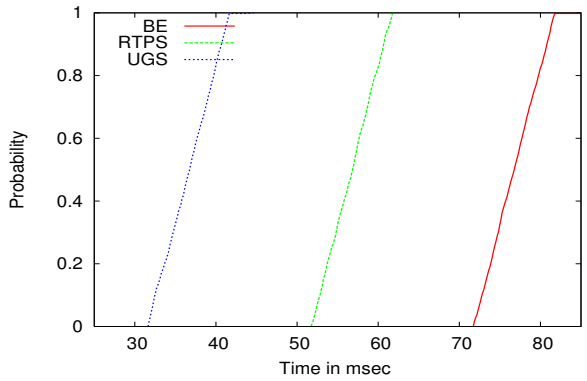

Fig. 2. Uplink probe: Observed round-trip latency

NPDU of 56 bytes, small enough to be transmitted in a single WiMAX subframe regardless of the modulation in use. Before sending the probe, the client stores a sequence number and the local time in the packet. When the server receives the packet, it adds its own local timestamp to the packet before echoing the packet back to the client. When the echoed probe is received by the client, the sequence number, the two time stamps, and the time at which the response was received are logged to a file. Packet send times are controlled by the *pthread_cond_timedwait()* facility, and the client may be configured to send randomly or synchronously.

The WiMAX network operates synchronously with respect to its 10 ms frame time. Nevertheless, even when probes are issued synchronously with an interprobe time that is an integral multiple of 10 ms, round-trip times experience a circular drift because the clocks of the WiMAX base station and the issuing host do not run at precisely the same rate.

For this reason, it is easier to understand the underlying dynamics when the probe packets are issued with random interprobe times. Our interprobe times were uniformly distributed in $[0.5, 1.5]$ seconds. With these random interprobe times, it is to be expected that observed round-trip times would be uniformly distributed across a time interval of 10 ms with the actual value of an individual observation being a function of the offset of the generation time within the 10 ms frame. For UGS, RTPS, and BE provisioning, 1200 probes were sent, requiring approximately 20 minutes of real time for each of the three experiments. No loss of probe packets was observed.

Figure 2 shows the empirical cumulative distribution of the measured round-trip times when 1200 probes were sent uplink from *wimax24* to *wimaxgw*. For each service class, the distribution is approximately uniform with a spread of 10ms as expected. The mean round-trip latencies of 36.69 ms (UGS), 56.74 (RTPS), and 76.77 (BE) are offset by 20 ms.

Being accustomed to seeing WiFi round-trip latencies consistently less than 5 ms, we were surprised to observe such large values on an otherwise idle WiMAX network.

In the remainder of this section we use empirical cumulative distribution functions to explore the underlying dynamics that produce these results.

3.1 Measuring one-way latency

Round-trip latency is easily measured to near microsecond accuracy using the Linux *gettimeofday()* facility. However, determining one-way latency is more challenging unless a global time source is available on both client and server, and in this case it was not. Nevertheless, the fact that client and server were connected both by WiMAX and a gigabit LAN (not shown in figure 1) make it possible to determine one-way latency to submillisecond accuracy even in the face of clock skew and clock drift.

Suppose measured round-trip latency over the gigabit LAN is t_{rg}. Then the best estimate for one-way delay is $t_{eg} = t_{rg}/2$. Now suppose t_{mg} is the measured one way delay obtained by subtracting the timestamp stored in the probe packet by the client sender from the timestamp stored by the server. This is the actual one-way delay if and only if the client and server clocks are synchronized. If clock skew exists, then t_{mg} is inaccurate by the amount of the clock skew, and thus the best estimate for present clock skew is $t_{eg} - t_{mg}$. Furthermore this estimate is bounded by $\pm t_{rg}/2$ because we know that the actual one-way time is greater than 0 and less than t_{rg}.

It should be clear that if we were trying to determine the one-way delay *over the gigabit LAN* this analysis is of no benefit. But since t_{rg} is on the order of 200 μs, we have an estimate of clock skew whose maximum error is on the order of $\pm$ 100 μs. Since both one way and round-trip times exceed 10 ms on the WiMAX path, using $t_{eg} - t_{mg}$ as the estimate of clock skew imparts an error of no more than 1% to measurements made on the WiMAX path. Therefore, we estimate the one-way delay on the much longer WiMAX path as $t_{1w} = t_{mw} + (t_{eg} - t_{mg})$ where t_{mw} is the difference between the time stamps in the probe that traversed the WiMAX network.

The above technique addresses the issue of instantaneous clock skew, but it does not address the issue of continuous clock drift. It was observed in the output of several runs that the clock on *wimaxgw* lost 0.1 seconds of time relative to the clock on *wimax24* in 1200 seconds. For the average one-way latency of 60 ms experienced by the BE probe, the amount of drift is thus $0.060 \times 0.1/1200 = 5\mu s$. Thus the clock drift *within a single probe period* is small relative to the known error of 100 μs. Nevertheless, it is not possible to ignore drift over a 1200 probe run of 20 minutes, and so clock skew must recomputed every probe period. This was accomplished by running concurrent probe processes on both paths, using a common pseudo-random number seed to keep them synchronized, and recomputing clock skew for every probe.

Using the above technique for computing outgoing latency, return path latency is computed by subtracting outgoing from round-trip. Round-trip and one-way latencies for each of the six experiments are shown in Table 3. The following characteristic behaviors may be noted in the table. Downlink latency is independent of scheduling type. Uplink latency increased by 20 ms from UGS to RTPS and from RTPS to BE. Uplink latency is longer than downlink latency for all scheduling types regardless of whether the client was uplink or downlink. Round-trip latency was increased by approximately 4 ms when the client was sending downlink. The scheduling dynamics that produce this behavior are described below.

TABLE 3
Observed Probe Latency

	UGS		RTPS		BE	
Uplink	μ	σ	μ	σ	μ	σ
Outgoing	19.86	2.91	39.81	2.85	59.87	2.88
Return	16.82	0.24	16.92	0.25	16.90	0.29
Round-trip	36.69	2.93	56.74	2.85	76.77	2.89
Dnlink	μ	σ	μ	σ	μ	σ
Outgoing	17.20	2.88	17.30	2.89	17.28	2.95
Return	23.04	0.23	43.04	0.32	63.01	0.29
Round-trip	40.24	2.88	60.34	2.95	80.29	2.93

The empirical distributions of the one-way and round-trip times provide insight. The one-way distributions of the outgoing uplink probe times are shown in Fig. 3. Note that the distributions have the expected uniform shape over a 10 ms interval with means of 19.86, 39.88, 59.87. Sample standard deviations of 2.91, 2.85, and 2.88 are consistent with the 2.89 standard deviation of the uniform distribution with spread 10.

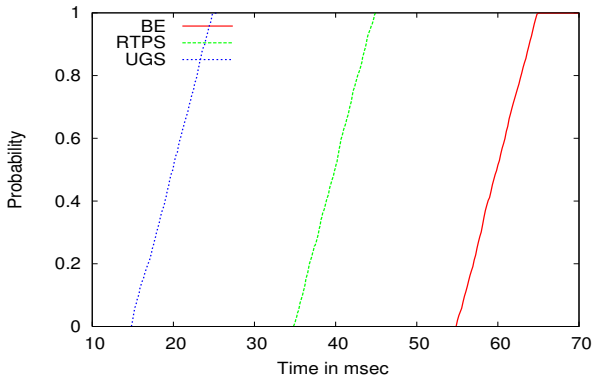

Fig. 3. Uplink probe: Observed outgoing latency

Unlike the uplink and round-trip distributions one would not expect to see a uniform distribution for the return latency of uplink probes. The packets arrive at the uplink host synchronized with the 10 ms frame time of the WiMAX system, and thus the strongly modal downlink distribution shown in Fig. 4 is to be expected. Here the observed means are 16.92, 16.90, and 16.82 ms.

To confirm that the the lengthy delays were indeed occurring within the WiMAX network we ran ICMP pings between *wimaxgw* and the base station and *wimax24* and the subscriber station. We obtained mean round-trip times of 322 μs and 775 μs respectively.

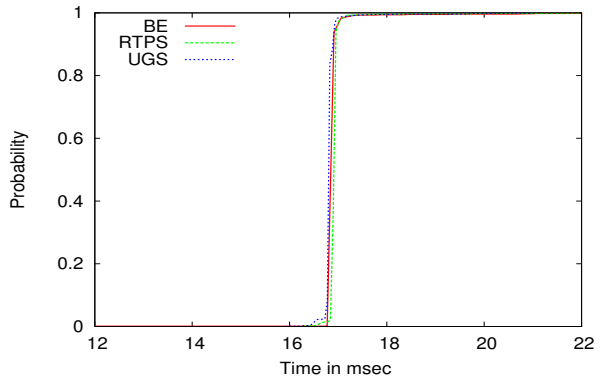

Fig. 4. Uplink probe: Observed return latency

3.2 WiMAX scheduling

Having obtained repeatable measurements of outgoing, return, and round-trip latency we sought to use them to infer the scheduling dynamics of the WiMAX network. We believe we were successful in doing so for both UGS and RTPS scheduling types. Nevertheless, some details of BE scheduling remain unclear. We first analyze the behavior of the uplink probes in which the client was located on *wimax24*.

3.2.1 UGS probes

Recall that the uplink latency of the probes was uniformly distributed between 15 and 25 ms. The underlying dynamics of this process are explained by Fig. 5. Each row of the table represents a single 10 ms frame that is equally apportioned between downlink (left) and uplink (right) components. The numerals in parentheses identify the approximate times that specfic events associated with the probe occur.

The probe packet is generated at some random time in frame 0 ms with mean offset at 5 ms. It arrives for transmission at the subscriber station in less than 1 ms at time (1).

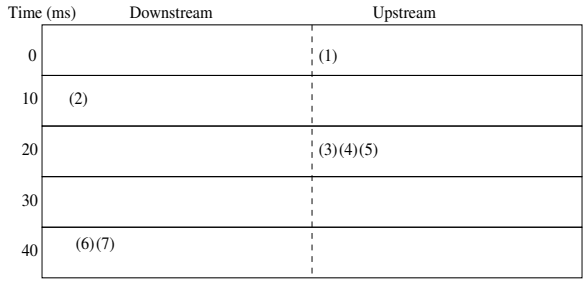

Fig. 5. Uplink probe: UGS timeline

In the next frame at time (2), a grant is received in the downlink MAP. Since there is no competing traffic, the grant is located near the start of the next uplink subframe, and the probe is transmitted uplink and received by the base station (3), forwarded to *wimaxgw* (4), and the response is received by the base station(5). Events (3)-(5) all occur in a timespan of less than 1 ms. The reply is finally received by the subscriber station at (6) and forwarded to *wimax24* at time (7) yielding the observed round-trip time of approximately 36.7 ms.

Two aspects of the scheduling algorithms can be inferred from this data. The subscriber station is provisioned to receive a grant in each UL-MAP. Thus, in frame 0 ms, it was granted the right to send in frame 10 ms. Nevertheless, it is apparent from the data that this grant can be used only for data that arrived *before the start* of the frame containing the grant. An analogous situation can be observed on the downlink side. Even though measurements indicate that the probe reply arrives at the base station in frame 20 ms, the response is never transmitted until frame 40 ms.

3.2.2 RTPS probes

Extension to RTPS is straightforward and illustrated in Fig. 6. Instead of receiving a grant at time (2) the subscriber station receives a poll granting it the opportunity to make a bandwidth request which occurs at time (3) in frame 20 ms. The grant is received at time (4) and the remainder of the exchange occurs as it does with UGS scheduling and yields the observed round-trip time of 57 ms.

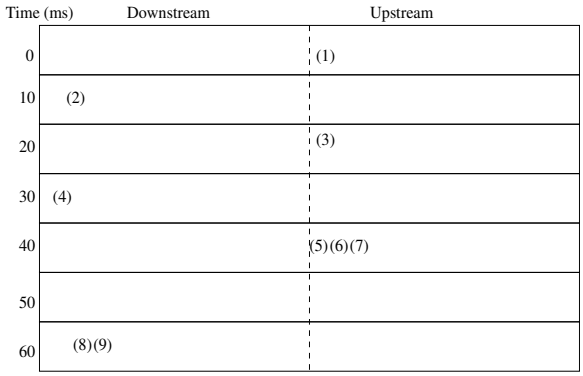

Fig. 6. Uplink probe: RTPS timeline

3.2.3 BE probes

Understanding of best effort dynamics is more difficult. In figure 7 points (1),(5), (6), (7), (8), and (9) are based upon observed one-way and round-trip latencies and represent the same events as their counterparts in the RTPS and UGS diagrams. However, the source of the extra 20 ms delay in the uplink as compared to the RTPS scenario is not clear. In the best effort timeline, point (2) represents receipt of an uplink MAP defining a contention opportunity for making a bandwidth request. Point (3) represents issuance of a contention request for bandwidth. This will never experience a collision on this otherwise idle network. Point (4) represents sending of a grant. Nevertheless, it is not possible to precisely identify the frames in which events (3) and (4) take place, and consequently the specific cause of the additional 20 ms delay is not clear.

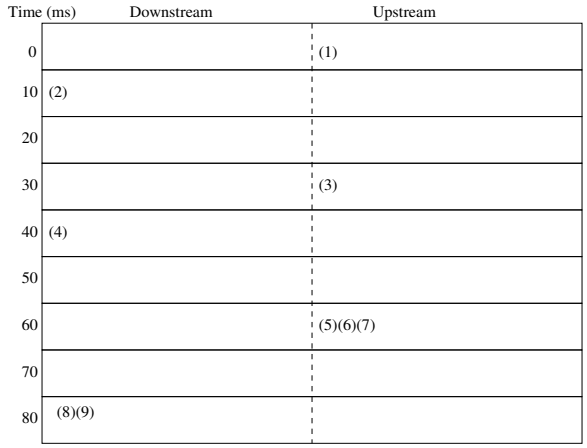

Fig. 7. Uplink probe: BE timeline

3.3 Downlink probes

The experiments were repeated with the client sending downlink from *wimaxgw* to *wimax24*. The data obtained supports the inferences derived from the analysis of the uplink probes. The distribution of downlink component times is shown in Fig. 8. The outgoing latencies are uniformly distributed and independent of the scheduling type. Return latencies, shown in Fig. 9 are now strongly modal and carry the characteristic 20 ms offsets that propagate into the round-trip latencies shown in Fig. 10.

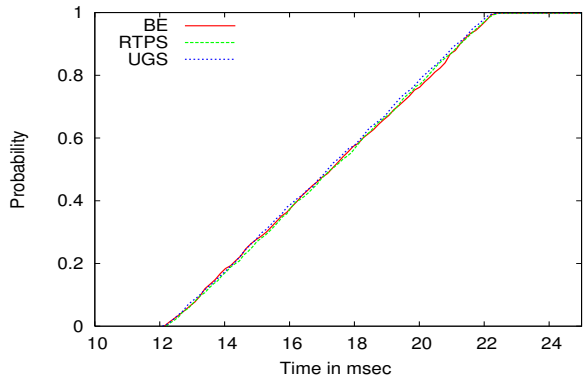

Fig. 8. Downlink probe: Observed outgoing latency

One anomalous aspect of the data is the unexpectedly high mean downlink latency which ranged from 17.2 ms to 17.3 ms. This effect can also be observed to a slightly lesser degree in the uplink probes. Given that the mean offset within frame of the time the transmission was initiated is 5 ms, a value closer to 15.5 ms would have been expected. Possible explanations include failure of the subscriber station to forward the probe in a timely manner or that a substantial amount of overhead preceded the probe in the downlink frame. The throughput study presented in the next section tends to confirm the substantial overhead hypothesis.

In summary, it can be stated that WiMAX provides inherently asymmetric round-trip latency, the latency is quite high for a LAN or MAN environment and thus definitely not well-suited for request-response applications with high throughput requirements.

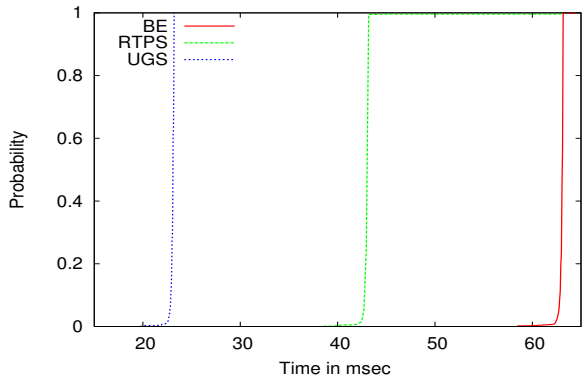

Fig. 9. Downlink Probe: Observed retrun latency

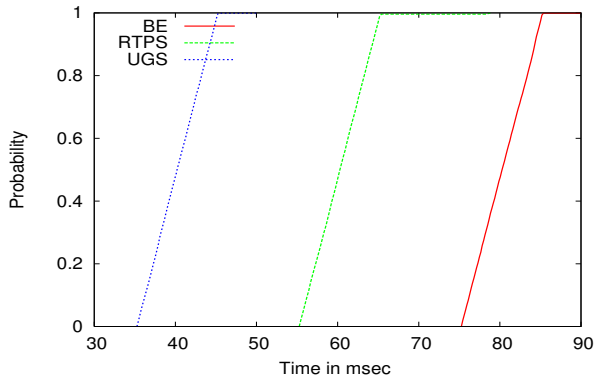

Fig. 10. Downlink probe: Observed round-trip latency

4 Throughput and TCP Dynamics

In this section we report upon throughput and TCP dynamics. As with the latency tests, these tests were conducted with three active subscriber stations. One of these was provisioned with a single RTPS flow. Thus the UL-MAP in each downlink subframe described a reservation slot in the next uplink subframe. However, no data was actually transferred on this flow. The subscriber station that carried the throughput tests was provisioned with a single best effort flow.

It was found that the scheduling algorithms on both the downlink and uplink properly constrained maximum throughput to precisely the provisioned amount when that amount was less than the capacity of the link. Therefore, the flows on which the throughput tests were conducted were configured with best effort scheduling and overprovisioned at 6 Mbps.

4.1 Throughput

All but two of the TCP throughput tests were conducted by serially running eight *iperf* transfers of 8 MB each between Linux hosts *wimax24* and *wimax01* shown in Fig. 1. During the throughput tests, the *tcpdump* utility

captured port-filtered raw packet traces at both ends of the connection. The traces were subsequently analyzed in a post-processing step.

The Airspan subscriber station that services *wimax24* is located in a third floor office window. It has a clear line of sight to the base station which is affixed to the roof of a nine story building at a horizontal distance of approximately 200 m. The office window is treated to resist thermal energy transfer, and this treatment produces almost a 20 dBm loss of signal power when compared to subscriber stations positioned outside at similar distances. The Airspan has a maximum power output of 20 dBm in comparison to 27 dBM at the base station. Because of these factors, the most efficient modulation achievable was 16-QAM 3/4 on the downlink and QPSK 3/4 on the uplink. Consequently, the 64-QAM 2/3 downlink and 16-QAM 1/2 uplink tests were performed using a M/A-COM subscriber station with 27dBM power output that was mounted in an automobile. These consisted of a single 8 MB transfer.

The modulation modes in use between the base station and a subscriber are not directly controllable by the system administrator. We controlled them by repositioning the subscriber station in the window sill and interposing journals of various thickness into the line of sight while monitoring the Web interface. It was observed that modulation schemes occasionally changed dynamically during an experiment yielding throughput numbers that were inconsistent. This situation was easily detected, and when it occurred, the entire experiment was repeated. In the absence of modulation changes throughput was very stable as shown in Table 4 which shows the elapsed time in seconds, downlink and uplink throughput, and network layer packet counts for the eight individual runs of the 16-QAM 3/4 downlink throughput test. Note that in seven of the eight runs an identical number of acknowledgments was transmitted.

TABLE 4
16-QAM 3/4 Downlink at Receiver

Elapsed time	Tx Kbps	Rx Kbps	Tx Pkts	Rx Pkts
17.670	68.841	3934.827	2924	5799
17.671	68.840	3934.818	2924	5799
17.662	68.875	3936.811	2924	5799
17.682	68.797	3932.341	2924	5799
17.662	68.899	3936.816	2925	5799
17.682	68.797	3932.352	2924	5799
17.671	68.840	3934.814	2924	5799
17.671	68.840	3934.818	2924	5799

The downlink throughput results are summarized in Table 5 and the uplink results in Table 6. The left column of these tables identifies the modulation being used. The next column shows the base station profile that was in use when the data was collected. The profile labeled $Default$ is the profile that was supplied with M/A-COM's UAS system. The profile labeled $Speedy$ contains changes suggested by M/A-COM to improve throughput. $Speedy$ specifies that the DCD and UCD

TABLE 5
Downlink IP Layer Throughput

Modulation	BS Profile	Kbps	Bits/Sym	Sym/Sec
QPSK 1/2	Default	1140	192	59.3
QPSK 3/4	Default	1705	288	59.2
16-QAM 1/2	Default	2268	384	59.1
BPSK	Speedy	662	96	69.0
QPSK 1/2	Speedy	1325	192	69.0
QPSK 3/4	Speedy	1985	288	68.9
16-QAM 1/2	Speedy	2641	384	68.7
16-QAM 3/4	Speedy	3936	576	68.3
64-QAM 2/3	Speedy	5256	768	68.4

TABLE 6
Uplink IP Layer Throughput

Modulation	BS Profile	Kbps	Bits/Sym	Sym/Sec
BPSK 1/2	Default	775	96	80.7
QPSK 1/2	Default	1553	96	80.9
BPSK 1/2	Speedy	848	96	88.3
QPSK 1/2	Speedy	1692	192	88.1
QPSK 3/4	Speedy	2540	288	88.2
16-QAM 1/2	Speedy	3361	384	87.5

be included in the downlink every other frame instead of every frame. On the uplink $Speedy$ provides six-symbol ranging opportunities every fifth frame instead of every frame and allocates two symbols instead of five for contention in each uplink subframe. The nominal savings is an average of 10 symbols per frame on the downlnk and 8.2 on the uplink.

The column labeled Kbps is the measured throughput at the IP layer in thousands of bits per second. It is followed by the number of data bits carried by each symbol. The column labeled Sym/Frame is the average number of physical layer symbols per frame consumed by the transfer as measured at the network layer. This value should be independent of the modulation and coding rate but dependent upon the base station profile and can be used to characterize the total PHY and MAC overhead.

The table shows that the $Default$ base station profile imposes an overhead of approximately 40% and the downlink but only 20% on the uplink. The $Speedy$ profile recovers 9 to 10 symbols per downlink subframe. Uplink overhead was reduced to approximately 12% using $Speedy$.

This data clearly indicates that a nominal 50/50 split provides an actual split that is strongly biased in favor of the uplink. A nominal 64/36 split should be used with the $Speedy$ profile to produce a approximately equal downlink and uplink capacity at the network layer.

4.2 Packet Loss

By sending UDP bursts of increasing size at 100 Mbps to both the base station and the subscriber station it was determined that both could buffer over 200 NPDUs of 1500 bytes each. TCP sender buffer size is limited to 128 KB in our Linux systems. Thus, a single TCP connection on an otherwise idle network can never

overflow the buffer space of either the base or subscriber station. Consequently, packets were never dropped for congestion during these tests.

The M/A-COM base station does not support ARQ. Therefore, highly aggressive strategies for dynamic modulation adjustment are not recommended and were not used. The results reported in tables 5 and 6 represent 98 total long-running TCP transfers, and the *tcpdump* traces showed that there was no packet retransmission in these runs.

Nevertheless, these results should be viewed as a best case scenario. They are likely to be achieved only when a fixed or nomadic subscriber station is stationary and has a direct line of sight to the base station that is unobstructed by moving traffic or other obstacles such as foliage. Even in a slowly moving vehicle, shadowing and multipath effects produce significant packet loss including complete loss of MAC layer connectivity. These effects will be fully discussed in our subsequent paper on coverage.

4.3 TCP Dynamics

The frame structure of the WiMAX system ensures that TCP dynamics will be somewhat bursty. In steady state flow, all packets that arrive in a single subframe are typically delivered with an interarrival time of approximately 180 μs. Then nearly 10 ms elapses before the next arrival. However, the limited capacity of the frame ensures that bursts never exceed 10 packets.

Figure 11 shows the empirical cumulative distribution functions of the interarrival times for eight downstream runs in which the sender was using 16-QAM 3/4 modulation. The resulting bit rate of 3936 Kbps, shown in table 5, corresponds to 3.2 IP packets of size 1500 bytes per WiMAX frame. Thus, if all frames are fully packed, it is to be expected that $1/3.2$ or 30% of the arrivals will experience a full frame delay while 70% will immediately follow their predecessor as shown in Fig. 11.

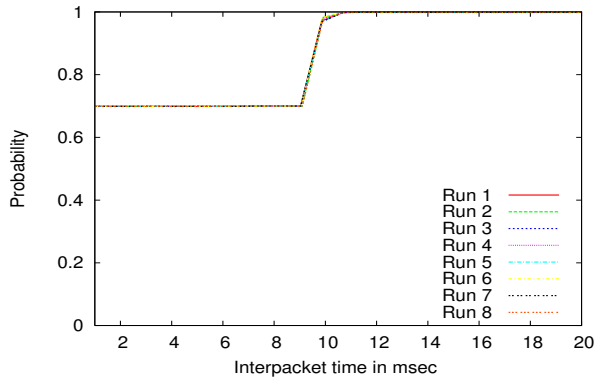

Fig. 11. Packet interarrival distribution

The interarrival distribution of the acknowledgment stream at the sender is shown in Fig. 12. It can been seen in Table 4 that this stream produced an average of 0.51 acknowledgments per packet. This corresponds to a acknowledgment rate of 1.66 acknowledgments per WiMAX frame. Therefore, in the most evenly distributed arrival process, 34% of the frames would carry one acknowledgment and the other 66% would carry two. In this best case distribution all of the first group and half of the second group for a total of 67% are preceded by a full frame delay and the other 33% immediately follow their predecessor. Figure 12 shows the actual distribution to be somewhat more bursty with 40% immediately following their predecessor.

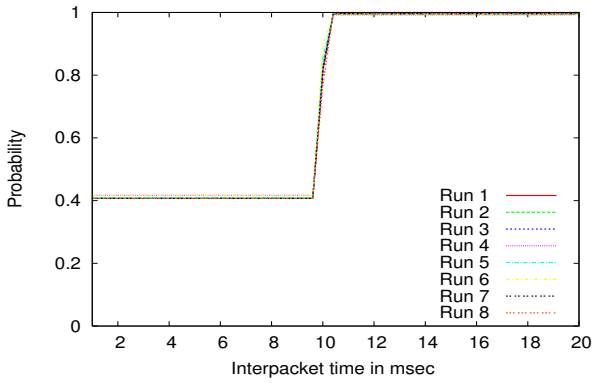

Fig. 12. Acknowledgment interarrival distribution

The actual distribution of acknowledgments within the 1641 frames of steady state operation of one of the eight TCP transfers is shown in Table 7. The first column carries the number of acknowledgments per frame that were observed. The second column has the number of frames observed to carry that number of acknowledgments. The aggregate column is the product of the first two columns. Thus 1641 frames carried a total of 2731 acknowledgments, and 1639 of the 2731 or 60% of them were the first acknowledgment in a frame. The other 40% followed a predecessor very closely as shown in Fig. 12. Nevertheless, it would not be fair to characterize this behavior as significant acknowledgment compression. Analogous interarrival distributions of packets and acknowledgments were observed in all throughput tests.

# of ACKS	Frames Having	% Frames Having	Aggregate ACKS
0	2	0.12	0
1	770	46.92	770
2	648	39.49	1296
3	219	13.35	657
4	2	0.12	8
Total	1641		2731

TABLE 7
Distribution of Acknowledgments

4.4 Source bursting and TSO

The packet and acknowledgment arrival streams, filtered through the leaky buckets of the base and subscriber

stations, are not particularly bursty. However, this is not the case with the initial sending of packets and consequently their arrival at the edge of WiMAX network. Furthermore, the presence of TCP segmentation offload (TSO) was seen to exacerbate this problem. A TSO capable network interface controller (NIC) can be passed an IP packet containing a TCP segment of up to nearly 64KB in size and will re-segment it into multiple IP packets of MTU size or smaller. TSO is *not the same* as IP fragmentation. Its operation can be observed in a *tcpdump* trace of the output stream through the presence of very large TCP payloads. The 17376 byte payload shown below constitutes 12 normal size payloads of 1448 bytes each.

```
16:27:17.072890 IP 192.168.90.24.53766 >
   192.168.90.1.33333: . 107985:125361(17376)
```

The objective of TSO is to reduce processor overhead at gigabit and higher speeds where 100,000 or more segments per second may be processed. However, even in the absence of TSO, two aspects of the Linux implementation of TCP induce bursty behavior at the sender: the implementation tries to do as much work as possible for a specific TCP connection in the context of the process that created the connection; and it also tries to avoid block/unblock "flapping" in which a process rapidly alternates between being blocked with full buffer quota and being unblocked with two available segments when an acknowledgment is received. Therefore, when a process becomes blocked due to full buffer quota, it will not be unblocked, and no further segments will be transmitted until a substantial fraction of the buffer space becomes free.

The following data taken from a *tcpdump* of the 16-QAM 3/4 downlink transfer at the sender is representative of steady state operation over the entire transfer. The sender, *wimax01*, does not support TSO. The left column encodes the operation as transmit or receive. The second column is the time since the start of the transfer in seconds. The next two columns specify the number of unacknowledged bytes and segments respectively. The last column is the usable window which is the offered window minus the number of unacknowledged segments. We see that the receive at time 17.025 reduces the number of unacknowledged segments to 44 and triggers a burst of 26 transmissions that all occur within the span of a millisecond. The next acknowledgment is received approximately 10 ms later at 17.035, and one or more acknowledgments continue to arrive every 10 ms until the number of unacknowledged segments drops below 45 again at time 17.105. This triggers another transmit burst of 26 segments.

```
R    17.025  63712     44   46
T    17.025  65160     45   45
T    17.025  66608     46   44
T    17.025  68056     47   43
       :
 --- 20 analogous records ---
       :
T    17.025  98464     68   22
T    17.025  99912     69   21
T    17.025 101360     70   20
R    17.035  98464     68   22
R    17.045  95568     66   24
       :
 --- 10 analogous records ---
       :
R    17.105  63712     44   46
T    17.105  65160     45   45
T    17.105  66608     46   44
```

The effect may be even more pronounced when TSO is enabled. On the uplink channel running in BPSK mode with network layer throughput at 848 Kbps, single segments as large as 65212 bytes (45 standard segments) were observed with *tcpdump*. With TSO enabled on the BPSK uplink channel, steady state behavior was periodic with a period length of 600 ms. Each period began with a a short transmit burst consisting of one or more TSO packets typically comprising a total of 44 segments of 1448 bytes each followed by a 600 ms period in which the 22 acknowledgments were more or less evenly distributed among the 60 frames. The resulting throughput was the maximum achievable on the otherwise idle channel.

Although this bursty behavior has no particular adverse effect when the network is otherwise idle, it would have strong negative consequences should a burst of 44 segments arrive at the subscriber station when buffer availability was very limited.

4.5 Managing TCP burstiness

The objectives of the mechanisms responsible for the introduction of source level burstiness are to improve throughput and reduce overhead at gigabit and higher speeds. At WiMAX speeds overhead reductions are negligible, and overall performance may be harmed. It is possible for the root user to disable TSO and to constrain the amount of write buffer space on a global basis, but this is not optimal for multi-homed hosts.

Another possible approach would be to modify Linux TCP so that it activates various optimizations on a per connection basis only in the presence of the traffic pattern for which the optimization was designed. This approach certainly adds complexity to a system that already has more than enough complexity.

At the application level it is possible to limit sender buffer space via the *setsockopt()* socket call[2]. Conflicting BSD and Linux semantics complicate this process, but via trial and error it was determined that when SO_SNDBUF was set to 8644 bytes[3] a WiMAX TCP connection could deliver the full throughput of 848 Kbps on a BPSK 1/2 channel. During steady state operation this configuration produced bursts of no more than five

2. In all of our systems we found that if TSO was enabled, the ability to set write buffer size with $setsockopt()$ was disabled.

3. Linux internally converts this to a buffer capacity of 17288, but this capacity limit includes storage required for *sk_buff* headers in addition to packet data.

packets and the number of unacknowledged packets did not exceed nine. Further reductions to sender buffer size produced loss of throughput, and more efficient modulation techniques require proportionately larger buffer capacity to be able to sustain full rate transmission.

4.6 Round-trip times

Round-trip times were measured for those transmissions that immediately followed a reception. Since the receipt of an acknowledgment completes the round-trip, it is to be expected for the round-trip time to be strongly modal with respect to the 10 ms frame time.

To obtain maximum throughput, it is necessary to maintain a population of data carrying packets at the base or subscriber station sufficient to ensure the full payload of each subframe is populated by data packets or fragments thereof. For uplink transfers, it is also necessary to ensure that the subscriber station has enough backlog that uplink capacity can be continuously allocated via the piggybacking process.

Additional sender buffer capacity beyond that which is necessary to support continuous transmission on the bottleneck air link cannot increase throughput. Beyond this point the magnitude of the round-trip time grows linearly according to Little's Law as a function of increasing packet population in the network.

Round-trip times experienced by the downlink 16-QAM 3/4 experiment are shown in Fig. 13. As previously described, the packets for which RTT measurements were taken were typically transmitted with a population of 45 unacknowledged packets in the system. Throughput is 328.37 packets per second, and so the approximate magnitude of the expected round-trip time is $46/328.57 = 0.140s$. Jitter in the system produces a second mode at 150 ms.

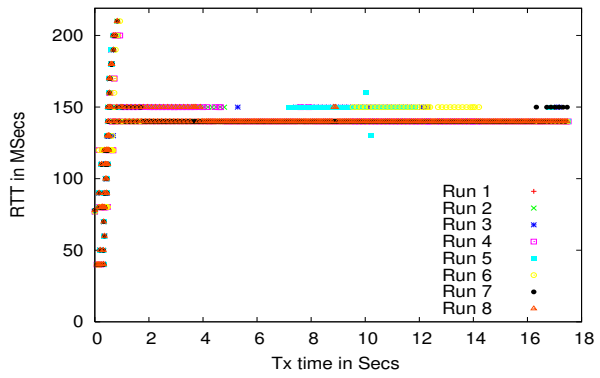

Fig. 13. Observed round-trip times

The uplink BPSK 1/2 transfer with SO_SNDBUF set to 8644 produced round-trip times with modes of 70, 80, and 90 ms.

4.7 Transport protocol and scheduling class effects

By running UDP $iperf$ tests with controlled bit rate, it was seen that maximum sustained throughput at the IP layer was the same as was achieved with TCP. TCP tests were also run with the scheduling class set to RTPS, and no change in maximum throughput was observed. The UGS scheduling class is not appropriate for unconstrained offered loads.

5 Related Work

Although there exists a considerable body of published research in the WiMAX domain, most of it is not directly related to our work. Development of specific scheduling algorithms and then using analytic or simulation models to evaluate their performance is the focus of several papers. The authors of [9] propose a scheduling algorithm for half-duplex subscriber stations operating in a network in which frequency division duplexing is used by the base station. They formally prove properties of the algorithm and then demonstrate via simulation its effectiveness in carrying a mix of VoIP and Web traffic.

The use of rate control on the packet arrival process for assuring the QoS guarantees for uplink RTPS or nRTPS traffic is studied in [10]. An analytic model is developed and validated via simulation. A RED-like mechanism is proposed for controlling that arrival rate of each uplink source that uses polled service. When the present polled service queue length is less than τ_{min}, the arrival rate is unconstrained. When greater than τ_{max} the arrival rate is constrained to some value λ_{min}. As the queue length grows from τ_{min} to τ_{max}, the maximum allowed arrival rate is continuously throttled until it reaches λ_{min} at τ_{max}. The technique is shown to stabilize delay in both steady state and transient conditions.

In [11] an ad hoc simulation developed by the authors is used to evaluate the capability of a simulated FDD network to provide differentiated services to video conferencing, VoIP, and data transfer workloads. Deficit round-robin (DRR) scheduling is used at the base station for downlink scheduling. Because DRR requires knowledge of the size of the head-of-line packet, weighted round-robin was used to schedule allocation of uplink bandwidth. A no-loss channel was also assumed. A related study, [12], uses simulation to evaluate the impact of using different physical layer frame sizes on both data and multimedia workloads. Other papers that report on the simulation of scheduling algorithms include [13], [14], and [15].

Evaluation of specific aspects of the MAC layer protocol itself is the focus of other papers. In [16] simulation is used in an investigation of the optimal number of contention slots in a WiMAX network. An OPNET simulation of the effectiveness of piggybacking compared to contention is presented in [17]. Simulation is used to demonstrate the impact of fragmentation and concatenation in a TDD WiMAX network in [18].

Simulation of the performance of extensions to the 802.16 standard is the theme of [19] in which it is proposed that dynamically varying priorities be associated with service flows. Other works use simulation or

analysis to model characteristics of the physical layer. Included in these are are [20] in which an OPNET simulation is used demonstrate the importance of dynamic modulation changes and [21] in which a physical layer model is used in a simulation study of coverage at 450 MHz and 3.5 GHz.

Aspects of both physical layer modeling and enhancing TCP performance are found in [22]. An OPNET simulation is used to evaluate the impact of ARQ retransmission delay on TCP performance. Because traffic flow on a TCP connection is typically asymmetric with full sized segments flowing in one direction and small acknowledgments flowing in the other, asymmetric adaptation of modulation is suggested. It is shown that overall throughput is optimized when a more robust but less spectrally efficient modulation technique is used on the acknowledgment channel than on the data channel.

Of the few studies that involve measurements taken on an operational WiMAX network, the most extensive is [23]. It is complementary to our work in that it uses a commercial network in which the authors have no control over the provisioning nor the level of competing traffic. The network in this study is characterized as "fixed" WiMAX and operated in Canada by two commercial service providers in the 2.496-2.699 GHz band. The authors' equipment was attached to the network via Motorola Expendience RSU-2510F subscriber stations. The providers limit downlink rates to 1.5 Mbps and uplink to 256 Kbps. The study compares the throughput and RTT obtained using four TCP variants with transfers in both downlink and uplink directions as a function of transmit buffer size. It is shown that a buffer size of 64KB is necessary to obtain near link speed downlink throughput with an RTT varying from 0.12 to 0.40 seconds. Throughput differences among the NewReno, Cubic, Vegas, and Veno TCP implementations was not significant, but Cubic TCP was shown to produce excessive retransmissions when using auto-tuned sender buffer space management.

A very brief paper [24] describes the performance of synthetically generated VoIP traffic over a real WiMAX network. The authors compare the performance of G.723.1 and G.729.2 CODECS as a function of the number of concurrent calls using the E-model as a metric. It is not clear if the network was carrying competing traffic at the time of the study.

Another study that employs an operational network is [25]. The focus of this paper is the development and evaluation of an end-to-end protocol for dynamic addition of service flows in a hybrid network that used 802.11e for access and WiMAX for back haul. The focus of this work is on the performance of dynamic service activation triggered by the subscriber stations. The authors worked with the equipment vendors in the design and implementation of the protocols.

6 Conclusion

In this paper we have described our experiences configuring and analyzing the performance of a WiMAX network. Our results are applicable to both fixed and nomadic subscriber stations operating at 4.9 GHz with output power limited to 27 dBm using low gain omnidirectional antennas under line of sight conditions.

Compared to WiFi technology, WiMAX provides significantly improved capability for provisioning QoS guarantees. Nevertheless, the capability is not without cost. Provisioning is a time consuming manual task. The system administrator must have a fundamental understanding of both PHY and MAC protocols to avoid unnecessary protocol overhead. Even in a fully compliant implementation, provisioning may be static in the sense that if a subscriber station that owns a UGS or RTPS service flow is powered on, then the flow consumes bandwidth regardless of whether or not the target application is even running.

Under uncongested conditions, a WiMAX network can be expected to impose considerably more latency than does WiFi. Nevertheless, as congestion grows, proper provisioning should make it possible to better manage the growth of latency than in even 802.11e WiFi nets. On our WiMAX testbed under uncongested conditions, downstream latency is independent of scheduling type. Upstream latency is higher than downstream for all type and strongly dependent on scheduling type. For all levels of congestion and all scheduling types any type of request/response application will see very poor throughput.

In contrast, any type of application using a pipelined protocol such as TCP can expect to obtain virtually 100% of the available capacity of an uncongested WiMAX network. The limited bandwidth of the WiMAX network produces reasonably small bandwidth-delay products even in the face of large latency.

Overhead imposed by the PHY and MAC layers was strongly asymmetric on our testbed. Uplink overhead of 12% is not significantly worse than cell header overhead in an ATM network, but a downlink overhead of 30% or more is clearly undesirable.

The long term future of WiMAX remains unclear. In many respects WiMAX versus WiFi today is the wireless counterpart to wired ATM LAN technology versus 100 Mbps Ethernet 10 years ago. At the time ATM LANs were faster and better suited to providing QoS guarantees, but Ethernets were much cheaper to deploy and easier to manage. Ethernet technology quickly evolved to gigabit speeds while still maintaining a cost advantage effectively relegating ATM LAN technology to a small niche role.

Acknowledgements

This project was supported by Grant No. 2006-IJ-CX-K035 awarded by the National Institute of Justice, Office of Justice Programs, US Department of Justice. The

project would not have been possible without the interest and cooperation of the Police Departments of Clemson University and the City of Clemson. The contributions of Sgt. Rick Clark of the Clemson University Police Department are worthy of particular note.

The authors also acknowledge the outstanding support provided by M/A-COM personnel in procuring and installing the equipment and in technical support. The contributions to the project of M/A-COM engineers Tom Brown and Mike Gaudette are especially noted.

Points of view in this document are those of the authors and do not necessarily represent the official position or policies of the US Department of Justice, the public safety agencies involved in the project, or M/A-COM, Inc.

REFERENCES

[1] "IEEE standard for local and metropolitan area networks part 16: Air interface for fixed broadband wireless access systems," *IEEE Std 802.162004 (Revision of IEEE Std 802.162001)*, pp. 1–857, 2004.

[2] "IEEE standard for local and metropolitan area networks part 16: Air interface for fixed and mobile broadband wireless access systems," *IEEE Std 802.16e2005 (Revision of IEEE Std 802.162004)*, 2005.

[3] J. Burbank and W. Kash, "IEEE 802.16 broadband wireless technology and its application to the military problem space," *Military Communications Conference, 2005. MILCOM 2005. IEEE*, vol. 3, pp. 1905–1911, 17-20 Oct. 2005.

[4] J. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX*. Upper Saddle River, NJ: Prentice-Hall, 2007.

[5] C. Eklund, R. Marks, K. Stanwood, and S. Wang, "Ieee standard 802.16: a technical overview of the wirelessMAN air interface for broadband wireless access," *Communications Magazine, IEEE*, vol. 40, no. 6, pp. 98–107, Jun 2002.

[6] A. Ghosh, D. Wolter, J. Andrews, and R. Chen, "Broadband wireless access with WiMax/802.16: Current performance benchmarks and future potential," *IEEE Communications Magazine*, vol. 43, no. 2, pp. 129–136, 2005.

[7] M. S. Kuran and T. Tugcu, "A survey on emerging broadband wireless access technologies," *Computer Netwworks*, vol. 51, no. 11, pp. 3013–3046, 2007.

[8] "FCC designates 4.9 ghz band for use in support of public safety," *FCC News*, February 2002. [Online]. Available: http://www.fcc.gov/Bureaus/Wireless/News_Releases/2002/nrwl0202.html

[9] A. Bacioccola, C. Cicconetti, A. Erta, L. Lenzini, and E. Mingozzi, "Bandwidth allocation with half-duplex stations in IEEE 802.16 wireless networks," *IEEE Transactions on Mobile Computing*, vol. 6, no. 12, pp. 1384–1397, 2007.

[10] D. Niyato and E. Hossain, "Queue-Aware Uplink Bandwidth Allocation and Rate Control for Polling Service in IEEE 802.16 Broadband Wireless Networks," *IEEE Transactions on Mobile Computing*, pp. 668–697, 2006.

[11] C. Cicconetti, L. Lenzini, E. Mingozzi, and C. Eklund, "Qality of service support in IEEE 802.16 networks," *IEEE Network*, vol. 20, no. 2, pp. 50–55, 2006.

[12] C. Cicconetti, A. Erta, L. Lenzini, and E. Mingozzi, "Performance evaluation of the IEEE 802.16 MAC for QoS support," *IEEE Transactions on Mobile Computing*, vol. 6, no. 1, pp. 26–38, 2007.

[13] R. Jayaparvathy, G. Sureshkumar, and P. Kanakasabapathy, "Performance evaluation of scheduling schemes for fixed broadband wireless access systems," *Networks, 2005. Jointly held with the 2005 IEEE 7th Malaysia International Conference on Communication., 2005 13th IEEE International Conference on*, vol. 2, pp. 6 pp.–, 16-18 Nov. 2005.

[14] A. Lera, A. Molinaro, and S. Pizzi, "Channel-aware scheduling for QoS and fairness provisioning in IEEE 802.16/WiMAX broadband wireless access systems," *Network, IEEE*, vol. 21, no. 5, pp. 34–41, Sept.-Oct. 2007.

[15] J.-C. Lin, C.-L. Chou, and C.-H. Liu, "Performance evaluation for scheduling algorithms in WiMAX network," *Advanced Information Networking and Applications - Workshops, 2008. AINAW 2008. 22nd International Conference on*, pp. 68–74, 25-28 March 2008.

[16] S.-M. Oh and J.-H. Kim, "The analysis of the optimal contention period for broadband wireless access network," in *Third IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOMW'05)*. Los Alamitos, CA, USA: IEEE Computer Society, 2005, pp. 215–219.

[17] R. Pries, D. Staehle, and D. Marsico, "Performance evaluation of piggyback requests in IEEE 802.16," *Vehicular Technology Conference, 2007. VTC-2007 Fall. 2007 IEEE 66th*, pp. 1892–1896, Sept. 30 2007-Oct. 3 2007.

[18] C. Hoymann, M. Putter, and I. Forkel, "Initial Performance Evaluation and Analysis of the global OFDM Metropolitan Area Network Standard IEEE 802.16 ," in *Proc. of European Wireless conference*, 2004.

[19] L. de Moraes and P. Maciel, "A variable priorities MAC protocol for broadband wireless access with improved channel utilization among stations," *Telecommunications Symposium, 2006 International*, pp. 398–403, 3-6 Sept. 2006.

[20] S. Ramachandran, C. Bostian, and S. Midkiff, "Performance Evaluation of IEEE 802.16 for Broadband Wireless Access," in *Proceedings of OPNETWORK*, 2002.

[21] T. Javornik, G. Kandus, A. Hrovat, and I. Ozimek, "Comparison of WiMAX coverage at 450mhz and 3.5ghz," *2006 International Conference on Software in Telecommunications and Computer Networks, 2006*, pp. 71–75, 2006.

[22] X. Yang, M. Venkatachalam, and S. Mohanty, "Exploiting the MAC layer flexibility of WiMAX to systematically enhance TCP performance," *Mobile WiMAX Symposium, 2007. IEEE*, pp. 60–65, March 2007.

[23] E. Halepovic, Q. Wu, C. Williamson, and M. Ghaderi, "TCP over WiMAX: A measurement study," in *Proceedings of MASCOTS 2008*, To Appear, 2008.

[24] N. Scalabrino, F. De Pellegrini, I. Chlamtac, A. Ghittino, and S. Pera, "Performance evaluation of a WiMAX testbed under VoIP traffic," in *Proceedings of the 1st international workshop on Wireless network testbeds, experimental evaluation & characterization*. ACM Press New York, NY, USA, 2006, pp. 97–98.

[25] P. Neves, S. Sargento, and R. L. Aguiar, "Support of real-time services over integrated 802.16 metropolitan and local area networks," in *11th IEEE Symposium on Computers and Communications (ISCC'06), 2006*. Los Alamitos, CA, USA: IEEE Computer Society, 2006, pp. 15–22.