An Experimental Investigation of VoIP and Video Streaming over Fixed WiMAX

Kostas Pentikousis, Jarno Pinola, Esa Piri, and Frerk Fitzek
VTT Technical Research Centre of Finland
Kaitoväylä 1, FI-90571 Oulu, FINLAND
Email: firstname.lastname@vtt.fi

Abstract—**Despite the significant interest in WiMAX technology and deployment, there are very few publicly reported measurements from testbeds and field trials. As such, most WiMAX studies employ simulation and modeling. This paper contributes to our understanding of what is realistically possible using off-the-shelf fixed WiMAX equipment today. We employ multiple competing traffic sources over a point-to-multipoint WiMAX topology and measure the capacity of the WiMAX equipment to handle a multitude of VoIP flows between subscriber stations while delivering a variable number of video streams. We measure throughput, packet loss, and one-way delay for both line-of-sight (LOS) and non-line-of-sight (NLOS) conditions. For the one-way delay measurements we synchronize the clocks of all testbed hosts with a software-only, open source implementation of the IEEE 1588 Precision Time Protocol. We compare these one-way delay measurements with those obtained when GPS-based synchronization is used.**

I. Introduction

Network practitioners are eager to see the results of the anticipated large public WiMAX rollouts in 2008, as they may change telecommunication markets significantly [1]. For example, in Finland, a country with low population density (16 people/km^2) but at the forefront of wireless communications developments, municipalities in the Kainuu Region venture jointly with the local network operator (KPO) to deploy fixed WiMAX in suburban and rural areas. The population density in the Kainuu Region is 4 people/km^2. Until recently, residents in such areas had to contemplate narrowband-only and expensive connectivity options, as the economics of covering them with other broadband technologies are prohibitive. Some claim that WiMAX will allow network operators to expand into less populated areas, increase market share, and even provide the means to narrow the so-called "digital divide".

The IEEE 802.16 family of standards (see [2], [3] and www.ieee802.org/16) has drawn the attention of network researchers and practitioners for several reasons. Networks based on these standards can provide point-to-multipoint last-mile broadband IP connectivity and play an important role in the evolution towards 4G. On the research side, some issues still need to be tackled at the lower end of the protocol stack, while the reference architecture proposed by the WiMAX Forum (see www.wimaxforum.org) has several aspects that need further definition. Optimizations are also needed in areas such as quality of service and mobility. Other issues of great interest for users, such as global roaming, are still unsettled.

In theory, WiMAX can deliver cell bitrates greater than 100 Mb/s, covering large areas (up to 50 km radius, from a single base station site using directional antennas), and serving tens of subscribers. Manufacturers have reported many successful trials and demonstrations with impressive throughput figures. Nevertheless, only a handful of empirical WiMAX studies employing commercial off-the-shelf (COTS) equipment have been published in the peer-reviewed literature. They all reported considerably lower capacities over point-to-point links. We are not aware of any empirical study that reports results from point-to-multipoint topologies. One plausible reason for this is that, despite the clear interest in the technology, WiMAX equipment is neither widely available nor inexpensive. In fact, currently, the WiMAX-related literature is based primarily on analysis and simulation, making use of the general properties of systems employing Orthogonal Frequency Division Multiple Access (OFDMA), and drawing on specifics from vendor-provided data, for the most part.

The first contribution of this paper is a thorough investigation of WiMAX point-to-multipoint performance in practice, thus filling a gap in the existing literature and forging a path towards more empirical work in this area. We employ multiple competing traffic sources over a point-to-multipoint topology with two subscriber stations (SS) and one base station (BS) and measure the capacity of our WiMAX equipment to handle a multitude of VoIP flows between SSs while delivering a variable number of video streams. We measure throughput, packet loss, and one-way delay for both line-of-sight (LOS) and non-line-of-sight (NLOS) conditions in our testbed. We expect our results to be of great value to researchers and practitioners alike. As we are interested in measuring one-way delay, the second contribution of the paper is a report on the effectiveness of a software-only, open source implementation [4], [5] of the IEEE 1588 Precision Time Protocol (PTP) [6].

This rest of this paper is organized as follows. In Section II we review and comment on related work. In Section III we present our testbed, explain our methodology and the use of PTP. It is widely anticipated that the next generation wireless networks will handle an exponentially larger amount of audio/visual (A/V) content than todays Internet. Due to their particular requirements, these two media types are put to the test in our lab. In Section IV we present our results and evaluate WiMAX as VoIP and video streaming backhaul in point-to-multipoint scenarios. We also contemplate on the use

WIOPT 2008, 1st – 3rd Apr 2008, Berlin, Germany.
 ISBN 978-963-9799-18-9
DOI 10.4108/ICST.WIOPT2008.3234

of PTP and compare it with GPS-based clock synchronization for WiMAX measurements. Finally, we discuss our results, lab experiences and challenges in Section V and conclude the paper in Section VI.

II. Related Work

Pioneering and closely related work to this was published by Scalabrino et al. [7], [8]. Using a fixed WiMAX testbed deployed in Turin, Italy, they focus on VoIP performance over WiMAX in particular when service differentiation is employed in the presence of significant amounts of elastic background traffic. Unfortunately, although their testbed included three SSs, the authors do not report any results from their simultaneous use. That is, their evaluation considers only point-to-point links. The same applies to the results reported by Grondalen et al. [9] from a fixed WiMAX field trial in Oslo, Norway. Their main means of evaluation are bulk TCP and UDP transfers. They measure throughput in both LOS and NLOS conditions and correlate it with received signal strength indicator (RSSI) values. Grondalen et al. find that their WiMAX system (employing the same modulation and FEC as this study) can deliver 9.6 Mb/s to a single flow in the downlink even at a distance of 5 km from the BS.

We have also measured 9.4 Mb/s in the downlink and 5.5 Mb/s in the uplink, using UDP bulk traffic with negligible ($<$0.1%) packet loss under direct LOS conditions. We recently obtained and reported these results [10] using the testbed employed in this study as well. Contrary to [9] our results are obtained with distances of only a few meters, but with considerably less transmission power (1 dBm vs. 28 dBm). In [10] we quantified the benefits of VoIP aggregation when fixed WiMAX is used to backhaul traffic. Other researchers have also explored the benefits of VoIP aggregation over WiMAX using simulation (see, for example, [11], [12] and the references therein). VoIP aggregation appears to be a powerful method to improve performance and increase capacity utilization, but is beyond the scope of this paper. Mignanti et al. [13] also report on FTP and VoIP performance over WiMAX in the Wind testbed in Ivrea, Italy. Their results indicate acceptable mean opinion scores for VoIP in a cell with a 2 km radius, but do not comment on overall (cumulative) throughput. Unfortunately, the results in [13] are not directly comparable to ours due to differences at the physical layer.

Due to the lack of extensive WiMAX-related public measurement reports, especially when considering point-to-multipoint WiMAX testbeds or field trials, we relate our work also with IEEE 802.11 measurement studies. Evidently, there are fundamental differences between IEEE 802.16 and 802.11. Nonetheless, the comparison between the two and the effects on VoIP (and video streaming) performance when multiple competing nodes claim resources from the BS can lead to interesting conclusions, as we will see. Of course, here we can only sample the literature on VoIP over 802.11 due to space considerations. Interested readers should also review, for example, [14] and the references therein.

Garg and Kappes [15] studied a scenario where duplex VoIP connections compete with a separate UDP flow in a single WLAN. They found that without competing VoIP traffic, UDP throughput was just over 6 Mb/s. For each, G.711-encoded VoIP connection added (10 ms voice frame interval and 132 octets packet size including RTP/UDP/IP headers), UDP throughput decreased by nearly 900 kb/s. When no competing UDP traffic was introduced, they found that a single 802.11b BS can handle in practice only six G.711 VoIP connections with an acceptable loss rate. Vasan and Shankar [16] also measured the effects of competing nodes on the aggregate throughput of an 802.11b BS. They show that the overall throughput decreases when the number of competing traffic sources, each sending maximum transmission unit-sized UDP/IP packets at maximum speed, increases. For example, a single UDP source can achieve 6.5 Mb/s, whereas ten competing sources can attain an overall (i.e. cumulative) throughput of only 2.8 Mb/s with a deviation of 1.5 Mb/s.

Dong et al. [17] show that 802.11b MAC mechanisms do not work well across all types of traffic scenarios and propose to dynamically adjust the duration of DCF and PCF modes in order to improve utilization. More recently, Marsh et al. [18] studied VoIP quality by measuring both application layer performance and MAC layer behavior in an IEEE 802.11b network. One test case considered the effects of bulk TCP competing traffic on the quality of G.711 VoIP over a single BS. They compared the round trip time (RTT) and jitter for zero to four competing TCP nodes sending at full rate in both managed and ad-hoc mode. In managed mode, even a single TCP source affects RTT and jitter significantly. In ad-hoc mode, RTT does not increase dramatically, but jitter does.

None of these results considers video streaming and VoIP, let alone WiMAX point-to-multipoint measurements. To the best of our knowledge, this is the first publicly available evaluation of video streaming over a WiMAX testbed considering the simultaneous use of two SSs in a single WiMAX cell that examines both VoIP and video streaming performance and reports accurate one-way-delay measurements.

III. Testbed and Methodology

Fig. 1 illustrates our experimental facility comprising an Airspan MicroMAX-SoC fixed WiMAX BS, operating at the 3.5 GHz frequency band, two subscriber stations (SS1 and SS2), and four PCs. SS1 is an Airspan EasyST and SS2 is an Airspan ProST. Symmetrically on the BS and SS sides, we connect GNU/Linux (kernel ver. 2.6.20-16, Ubuntu 7.04) PCs with Gigabit Ethernet PCI cards to act as traffic sources/sinks. We measure performance under both LOS and NLOS conditions.

A. LOS and NLOS Measurements

In the LOS measurements, the distance between BS and SS1 and SS2 was 5 and 10 m, respectively. The measured RSSI values were -48 dBm at both SSs. All tests were performed in our laboratory, where conditions are static, even though there can always be a degree of variance on a wireless link.

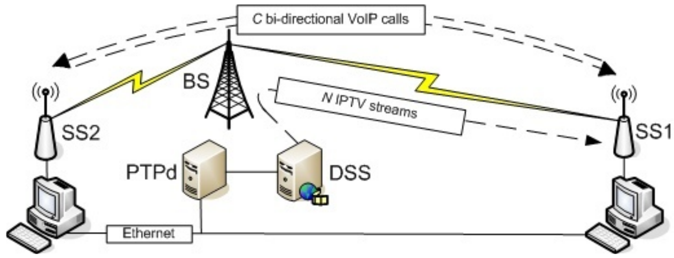

Fig. 1. Schematic of our WiMAX testbed.

TABLE I
TESTBED CONFIGURATION

Base station	Airspan MicroMAX-SoC
Subscriber stations	EasyST (SS1) & ProST (SS2)
Frequency band	3.5 GHz
Channel bandwidth	3.5 MHz
PHY	WiMAX 16d, 256 OFDM FDD
BS and SS Tx power	1.0 dBm
MAC scheduling	Best Effort
LOS Measurements for SS1 and SS2	
Uplink & downlink modulation	64 QAM (FEC: 3/4)
BS-SS1 distance	5 m
BS-SS2 distance	10 m
NLOS Measurements for SS1	
BS-SS1 uplink & downlink modulation	16 QAM (FEC: 3/4)
BS-SS1 distance	15 m

For this reason, we monitored the WiMAX equipment to ensure that the key parameters remained unchanged during the entire duration of the tests. We noted that despite the short distance and direct LOS, the default "adaptive" modulation scheme employed by SS1 continuously alternated between 64 QAM 3/4 and 64 QAM 2/3. This alternation made the measurement results to fluctuate noticeably, so we opted to disable the default adaptive modulation scheme at SS1 and fix the modulation to 64 QAM with 3/4 FEC. We did not observe similar problems with SS2.

In the NLOS measurements, the path between the BS and SS1 was blocked by a solid steel reinforced concrete wall; the BS-SS1 distance was 15 m. The measured RSSI at SS1 was -63 dBm. SS2 was fixed in the same location as in the LOS measurements. In the NLOS tests, SS1 was configured to use adaptive modulation. Once SS1 synchronized with the BS, its modulation was set to 16 QAM with 3/4 FEC throughout the entire set of NLOS measurements. The modulation at SS2 and its RSSI was the same as in the LOS measurements. Table I summarizes our testbed configuration.

B. Traffic Generation

In order to emulate a set of IPTV streams, we captured 20 minutes of live IPTV unicast transmission and created a packet trace. The captured video stream was in H.264/AVC format (also known as MPEG-4 Part 10) [19] and the accompanying audio stream was encoded in MPEG-1 Audio Layer II (also known as MP2) [20]. The IPTV content was streamed using the Darwin Streaming Server (DSS), an open source RTP/RTSP [21]–[23] server. We chose to emulate a music video TV channel and thus configured DSS to stream the video at 512 kb/s (360×288, 25 f/s) and the audio at 192 kb/s, emphasizing audio over video quality. We connected DSS with a receiving host using a Gigabit Ethernet switch and collected the video stream packet trace at the receiver side using Wireshark. We recorded very low delay and delay variance, no RTP packet loss, and no RTSP message exchanges.

The captured video stream, often with dramatic changes in scenery, visual effects and so on, has a variable bit rate (VBR). The total packet sizes of the video varied greatly, with the major mode at 1492 bytes. Based on the Wireshark trace, we created a new trace file with all packet sizes and interarrival times. All packet traces are available for free from the authors. For each of the independent replications reported below, we "play back" N IPTV A/V streams starting from a random point in the 20-minute long IPTV packet trace, and wrap around if needed. Each run lasts 60 s. We use JTG [24] to generate the trace-driven IPTV streaming traffic. JTG is a simple, flexible, and configurable open source traffic generator, which can be used in a command-line fashion in GNU/Linux.

Based on the Wireshark trace, we emulated the corresponding IPTV audio stream using constant bit rate (CBR) traffic with the total packet size fixed at 634 bytes (including codec payload and RTP/UDP/IP headers). The video and audio parts of the IPTV traffic are separated by DSS and streamed to different ports. This separation allows us to study in a straightforward manner the relative performance of VoIP and IPTV A/V over a congested fixed WiMAX link. The source of the N A/V streams is located in the wired part of the topology in Fig. 1, while the sinks are in the domain of SS1.

In addition to the N A/V streams, we injected C synthetic duplex, bidirectional VoIP flows with source/sink pairs in the domains of SS1 and SS2 as shown in Fig. 1 in all of our tests, again using JTG. We chose to experiment with Speex [25], an open source audio codec specially designed for VoIP applications over packet switching networks. Speex is designed to be robust against packet loss and has been incorporated in several applications, including Microsoft's Xbox Live. We emulated C Speex VoIP flows each with a wideband codec bitrate of 12.8 kb/s using JTG. For each VoIP flow, JTG generates 50 packets/s with 32 bytes of codec payload, thus leading to an effective application bitrate of 17.6 kb/s (including RTP headers). After adding a total of 28 bytes of UDP and IP headers, each JTG instance injects 28.8 kb/s of total emulated Speex CBR traffic into the network.

C. Host Clock Synchronization

For high-precision one-way delay measurements, accurate clock synchronization is necessary, taking care of both absolute time and clock drift at different hosts in the network. When only round-trip delay measurements are performed, the critical aspect is clock drift only. Lack of accuracy in the absolute time is not harmful. However for the one-way delay measurements we are interested in this paper, both absolute time and clock drift are important. Before starting the full range of experiments, we used QoSMET [26] on two Windows PCs with GPS clocks (Trimble's Acutime 2000 Synchronization Kit) to measure the one-way delay on the BS-SS2 WiMAX link. QoSMeT is a highly-accurate proprietary measurement tool developed by VTT. According to QoSMeT, $\tilde{d}_{downlink} = 8.7$ ms ($\tilde{d}$ denotes the one-way delay median) and $\tilde{d}_{uplink} = 23.5$ ms.

Based on these measurements, the use of Network Time Protocol (NTP)-based synchronization [27] was ruled out, as the one-way delay in our fixed WiMAX testbed is in the same order of magnitude as the accuracy of NTP (approx. tens of ms). Moreover, since we possess only two GPS clocks, we cannot perform measurements with multiple VoIP and video clients/servers. Thus we opted to use an IEEE 1588 Precision Time Protocol (PTP) [6] open source server (PTPd [5]) to synchronize the clocks at all hosts. Similarly with NTP, PTP synchronization messages are sent over the network.

Although PTP injects a very small amount of traffic when compared with the rest of the sources in our tests, it is preferable that PTP signaling does not interfere with the measured traffic. Thus we employed *two* network interface cards at each PC and routed the clock synchronization signaling and the VoIP and A/V streams separately, as illustrated in Fig. 1. We monitored both interfaces at each node with Wireshark and verified that this was the case indeed. In short, each testbed PC was connected to the PTP server via a dedicated Ethernet network interface and the measured traffic flows had no effect on synchronization and vice versa. Based on [28], this means that host clock synchronization within 100 μs is feasible.

IV. Results

Our baseline tests determined that our WiMAX testbed can sustain 100 simultaneous one-way Speex streams in the uplink, irrespective of whether one or two SSs were connected to the BS. In order to test VoIP backhauling inside the same WiMAX cell we introduced $C = 50$ simultaneous, bidirectional flows, as explained in §III-B, yielding an application goodput (Speex payload plus RTP header) of 880 kb/s. This is only 16% of the maximum uplink goodput of 5.5 Mb/s, measured with MTU-sized UDP packets. This cannot be entirely explained by the large UDP/IP header overhead (87.5%), and indicates that the tested WiMAX equipment does not handle (a large amount of) small IP packets as competently as MTU-sized packets.

The measured maximum downlink capacity with MTU-sized UDP packets was 9.4 Mb/s, which clearly exceeds the uplink capacity and leaves plenty of bandwidth for downloading for the subscribers inside the WiMAX cell. As we are interested to measure the performance of the tested equipment at its capacity limit in a scenario where VoIP and A/V streaming take center stage, we emulated a video on demand service to be used by some of the subscribers in parallel with the VoIP conversations. By gradually increasing N, the number of synthetic IPTV A/V streams served via the BS to SS1, we determined the "breakpoint" of the downlink. For each N, we repeated the run ten times. Our measurements are presented in the following subsections in "box-whisker-plots",

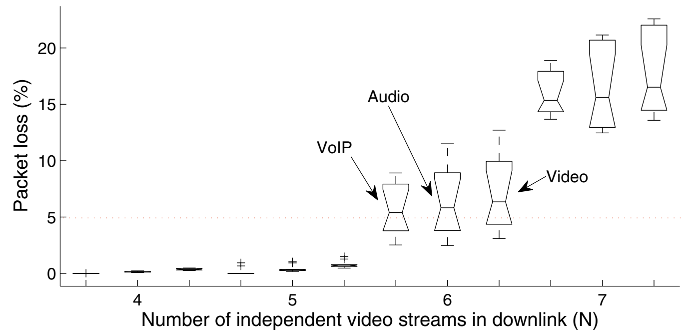

Fig. 2. Measured packet loss at SS1 under LOS conditions.

or simply boxplots. The box in each figure contains the middle 50% of the measured values. The line in the middle represents the median, the top and bottom of the box correspond to $Q3$ and $Q1$, respectively. Values outside the whisker lines, shown as crosses, are considered outliers.

A. Packet Loss

The measured packet loss rates, l, at SS1 for the video, audio and VoIP are depicted in Fig. 2. The box plots show the ten averages of packet loss for the VoIP and A/V streams, as recorded by JTG. The BS-SS1 WiMAX downlink can handle $N \leq 5$ simultaneous A/V streams in parallel with the VoIP traffic with negligible l. When $N = 6$, l increases rapidly, and $\tilde{l} > 5\%$, which is unacceptable. Even for the Speex codec, which is the most robust and tolerant to packet loss of the three codecs emulated, this level of losses would be degrading performance considerably. Nevertheless, the $C = 50$ VoIP flows seem to suffer fewer drops on average: ($\tilde{l}_{VoIP} < \tilde{l}_{audio} < \tilde{l}_{video}$). The IPTV video streams suffer packet losses that exceed 6%, for both video and audio, which could not be handled (i.e. concealed) satisfactorily by a real-world IPTV client. When $N > 6$ packet losses exceed 15%, at which point all receivers observe abysmal performance.

If we turn our attention to the performance of the 50 biderectional VoIP flows at SS2, we note a trivial amount of packet loss (not shown due to space restrictions). During the measurements, the highest l observed at the SS2-BS uplink was 0.4%, which is negligible for Speex.

B. Application Throughput

Figures 3-5 illustrate the application throughput or goodput, g, for the three types of traffic over the fixed WiMAX link. The difference between the VBR H.264/AVC encoded video streams and CBR audio and VoIP is noticeable. The fluctuation of the video stream throughput is more extensive. This is in part due to the variability of the wireless transmission, but also due to the random starting points in the replayed video packet trace. Note that due to the random starting points in the video trace, g_{video} ranges between 490 kb/s and 512 kb/s, when $N \leq 5$, as one would expect. Meanwhile, the throughput of the audio and VoIP streams remains very close to 178 kb/s and 17.6 kb/s, respectively. Note that although DSS was configured

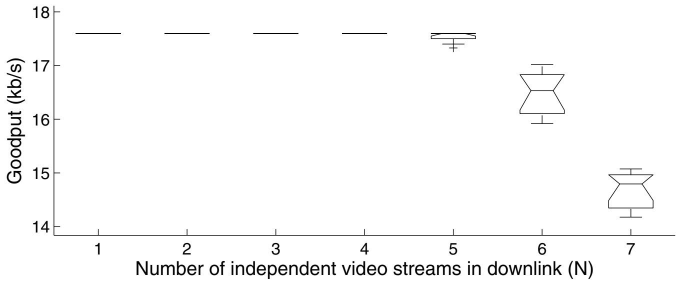

Fig. 3. Measured VoIP goodput at SS1 under LOS conditions.

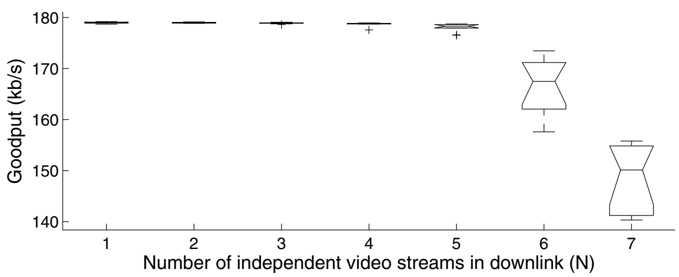

Fig. 4. Measured audio goodput at SS1 under LOS conditions.

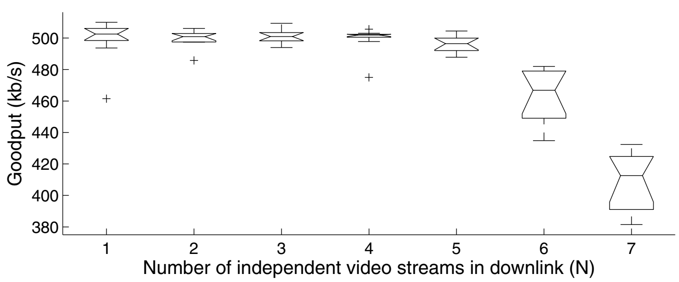

Fig. 5. Measured video goodput at SS1 under LOS conditions.

to transmit the audio at 192 kb/s, the captured audio packet trace recorded a somehow smaller bitrate.

When $N > 5$, the capacity of the WiMAX downlink becomes a restrictive factor and the median goodput of all traffic types starts to fall (see Figures 3-5). The spread of the average goodput in different runs starts to increase, as packets are dropped due to backlogs at the WiMAX interface. When the normalized values of g are examined, we note that the behavior of the three different traffic types is practically the same when $N > 5$, which is the "breakpoint" of the WiMAX downlink for this scenario.

The measurement data from SS2 (not depicted), on the other hand, showed that VoIP goodput is unaffected when the entire downlink capacity of the BS is used. VoIP goodput ranges between 17.5-17.6 kb/s even when $N > 5$. This is a significant difference with previous empirical studies that employed IEEE 802.11 access networks. In particular for VoIP and A/V streaming, and recalling [15], IEEE 802.16 seems to have an advantage over 802.11, as the relative uplink and downlink state of high resource utilization does not affect open-loop traffic flows. Of course, the tested WiMAX equipment does not have any mechanisms to support dynamic allocation of more bandwidth to the downlink at the expense of the uplink.

C. PTPd vs. GPS-assisted Clock Synchronization

As mentioned above, we employed PTPd [5] to synchronize all PCs in our testbed, in order to be able to accurately calculate one-way delay. In principle, with GPS-based synchronization, COTS PCs can achieve sub-μs accuracy with respect to UTC. In practice, on a Windows-based system this may actually mean an accuracy of 50 μs at worst.

In order to compare the accuracy of GPS- and PTP-based synchronization, we emulated part of QoSMeT's operation using JTG. Note that QoSMeT has many more functionalities, allowing multipoint measurement as explained in [26]. A bidirectional data stream with a bitrate of 64 kb/s and packet payload size of 240 bytes was sent over the otherwise unused BS-SS2 WiMAX link for a total of 25 minutes. The characteristics of this data stream are comparable to the G.711 VoIP actual codec which is used to create the packet stream in the GPS-based measurements performed with QoSMeT. For the PTP measurements, the delay and jitter characteristics of the packets were measured and recorded by JTG. Of course, we do not claim that the QoSMeT- and JTG-generated data streams are identical.

Figures 6 and 7 illustrate that PTP achieves first-rate synchronization accuracy in our testbed. In particular for the downlink (BS-SS2), see Fig. 6, PTP allows us to measure one-way delay with accuracy similar to that when GPS synchronization is used. On the uplink SS2-BS, however, one-way delay is not as accurately measured when employing PTP as when using GPS clocks. As can be seen from the inset boxplots of Fig. 7, when we use PTPd the measured median one-way delay, $\tilde{d}$, appears to be larger than it is in reality.

The measurements using QoSMET show that the median one-way delay in our fixed WiMAX testbed is 8.7 ms (downlink) and 23.5 ms (uplink). With all hosts synchronized within hundreds of μs, this allows for measurement granularity of one to two orders of magnitude larger than the one-way delay. At this stage, and wishing to avoid having to deploy (expensive) GPS-clocks on all hosts in the testbed, we are satisfied with the synchronization accuracy provided by the software-only implementation of PTP. In fact, we are quite confident about the one-way delay values reported in the following subsection and claim that the downlink one-way delay measurements are as good as if we had used GPS clocks at all hosts for this set of tests. With respect to the uplink one-way delay measurements, we can say that the results with PTP synchronization are slightly pessimistic on median.

D. One-way Packet Delay

The one-way packet delays, d, in the BS-SS1 downlink, as measured by the packet interarrival times at SS1, are quite similar across all traffic types, as shown in Fig. 8. Overall, $\tilde{d} < 150$ ms when $N \leq 5$. This range of one-way delays can be tolerated by all applications involved in the examined scenario. When $N = 6$, the median value of the interarrival times jumps to over 600 ms for all traffic types, as the majority

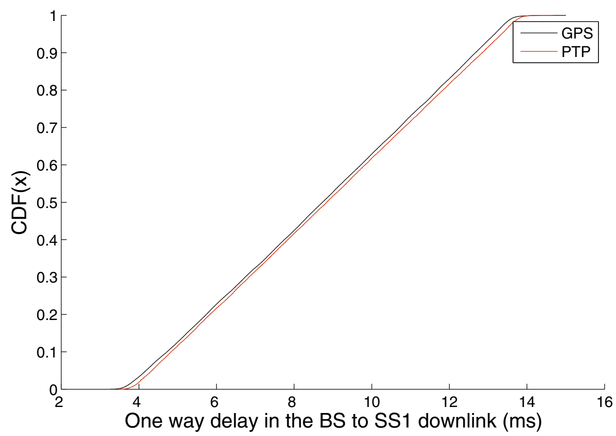

Fig. 6. Comparison between GPS- and PTP-based one-way delay measurements for the BS-SS2 downlink.

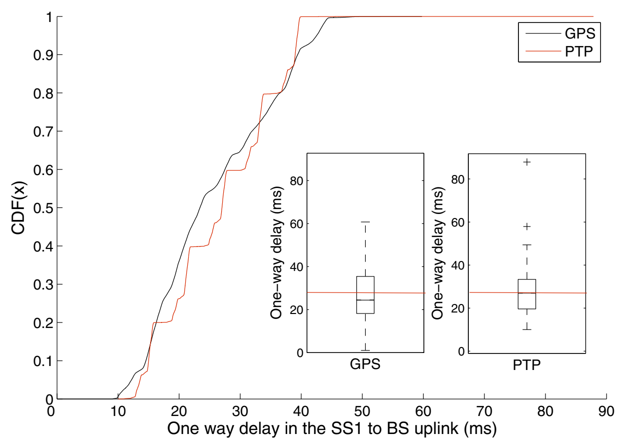

Fig. 7. Comparison between GPS- and PTP-based one-way delay measurements for the SS2-BS uplink.

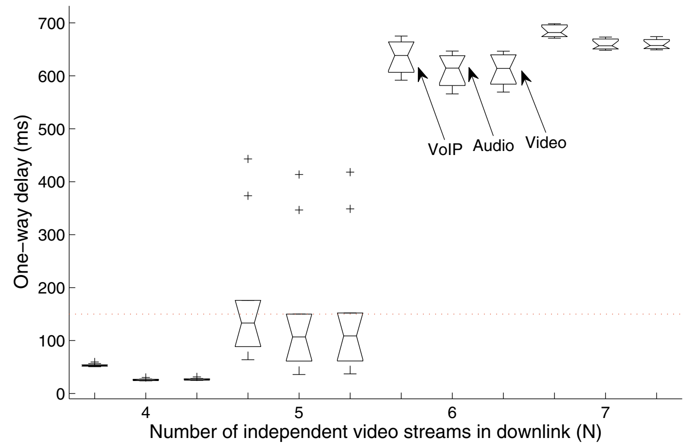

Fig. 8. Measured one-way delay at SS1 under LOS conditions.

of the received packets have queued through full buffers in the network interfaces of the test system. This range of one-way delays can be handled with adequate buffering for the IPTV streams, but not for the VoIP calls.

To sum up, under LOS conditions our testbed can sustain $C = 50$ emulated bidirectional Speex-encoded VoIP calls within the same WiMAX cell and $N = 5$ simultaneous emulated IPTV streams with negligible packet loss, adequate application-level throughput, and one-way delays within proper bounds. Recall that our BS does not employ any QoS mechanisms; these results are based on the "best effort" fixed WiMAX profile.

E. NLOS Measurements

The NLOS measurements are consistent with the results presented above under LOS conditions, regarding application throughput, packet loss and one-way packet delay. Note that during the NLOS measurements, SS2 uses the higher modulation scheme (64 QAM). Of course, given the different modulation scheme (16 QAM) employed automatically to cope with the weaker signal strength, the overall BS-SS1 link capacity is smaller. Because of this, performance deterioration occurs with a smaller N, it is more rapid, and more varied than in LOS conditions. We found that the BS-SS1 WiMAX downlink is capable of delivering only one IPTV A/V stream in addition to the $C = 50$ ongoing VoIP conversations. With $N = 2$, packet losses and one-way packet delays for the A/V are very high, 12% and 500 ms on median, respectively (figures omitted due to space restrictions). Under such conditions the received quality of the IPTV stream would be simply unusable. The same conclusion applies to the VoIP traffic as well.

On the other hand, examination of the traces at SS2 showed that the quality of the VoIP flows sent from SS1 to SS2 via the BS under NLOS conditions were received without serious problems at SS2. As Fig. 9 illustrates, goodput ranges between 17.4 and 17.6 kb/s, while the highest measured packet loss rates is just over 1%. The one-way packet delay does not exceed 80 ms, even when $N = 3$.

V. Discussion and Future Work

Grondalen et al. [9], after correlating RSSI and throughput, found that in practice high throughput levels can be attained at distances up to 5 km from the BS, indicating that the results of Section IV might be applicable, to some extent, to outdoor environments as well. Of course, this remains to be verified in future work, but note that our BS transmits at only a fraction of its maximum capacity (1 dBm when the max is 32 dBm). Our experiments were performed in a laboratory environment, which enables more reliable and stable conditions for capacity measurements than outdoor field trial. Nevertheless, wireless technologies, even in laboratory conditions, are prone to disturbance and therefore result variations are to be expected.

The PHY layer properties are very important even though the actual measurements only take place at or above the MAC layer. This is in particular the case for VoIP which typically generates streams with small packet sizes. The wireless modulation scheme has to be monitored too. Often devices are set to use adaptive modulation, which may interfere with the measurements. As mentioned earlier, we fixed the modulation at SS1 to 64 QAM (3/4 FEC) for the LOS measurements, in

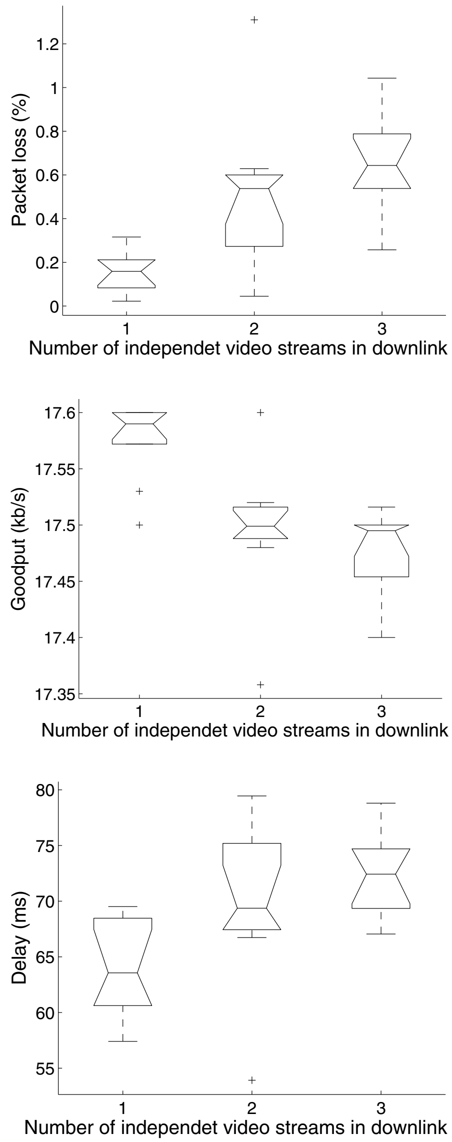

Fig. 9. Measurement results for VoIP under NLOS conditions at SS1

order to avoid having to deal with the perhaps buggy implementation of the adaptive modulation algorithm. In WiMAX testbed measurements, just like in other wireless testbeds, such aspects need to be addressed explicitly and reported candidly in addition to other relevant software and network layer issues. Our testbed consisted of IEEE 802.16d equipment following a subset of the physical layer specifications specified in [2]. In the future, these measurements will be repeated using 802.16-2005 [3] equipment, addressing mobility scenarios as well.

The results presented in Section IV are interesting from a different perspective as well. Although the actual goodput for VoIP is quite low when compared to the maximum attainable uplink goodput (Table II), we find that WiMAX is promising for backhauling multimedia traffic, especially if we consider the results reported in related empirical studies of VoIP over IEEE 802.11. When the performance of our WiMAX testbed is compared with the performance of a IEEE 802.11g access point (with nominal capacity of 54 Mb/s), WiMAX is considerably more efficient than WiFi. Our preliminary measurements with IEEE 802.11 and emulated, bidirectional Speex VoIP "calls", indicate that our WiMAX equipment can sustain twice as many flows as WiFi does within the same cell. Of course, this reflects fundamental differences at the PHY and MAC layers of IEEE 802.11 and 802.16. Nevertheless, the price differential between currently available WiMAX and WiFi equipment is very large and, let us not forget that, in addition, WiMAX operates in licensed bands whereas WiFi operates using unlicensed spectrum.

TABLE II
WiMAX Median Link Utilization

Median Utilization	BS-SS1 $N = 5$ & $C = 50$	BS-SS2 $C = 50$	BS-SS2 max UDP goodput
Downlink	45.3 %	9.4 %	9.4 Mb/s
Uplink	16 %	16 %	5.5
Aggregate	34.5 %	11.8 %	14.9

One the other hand, one could argue that the scenarios investigated indicate that WiMAX is not particularly efficient with this kind of traffic when compared with other wired local loop alternatives. Moreover, QoS in WiMAX is a much touted feature, yet experience from previous ambitious technologies, including ATM, casts doubt about the use of sophisticated features in today's Internet. When considering WiMAX without QoS support, such as the one investigated in this paper, we could argue that using WiMAX as backhaul rather than as network access technology may be more promising. When WiMAX is used as a point-to-point "bit pipe", various methods, including multiplexing/aggregation and header compression, can be used to increase efficiency.

VoIP performance over cellular networks has also received significant attention. Recently, for example, Wang et al. [29] evaluated VoIP over HSDPA using simulation. Each VoIP flow had, on average, 38-octet packets, including RTP/UDP/IP headers. They study the benefits from employing Robust Header Compression (ROHC) [30], which compresses all headers to 3 octets only. Thus, each VoIP flow has a target rate of 15.2 kb/s. Wang et al. show that with a 100 ms delay budget for VoIP, the HSDPA cell capacity is 86 flows where >95% of all packets are received correctly within the specified delay budget. Studying the potential of ROHC for WiMAX is an item in our current research agenda.

VI. Conclusion

Until recently, nearly all studies of WiMAX technology were performed using simulation and modeling, due to the lack of available and inexpensive equipment. In this paper, we investigated WiMAX point-to-multipoint performance in practice, filling a gap in the existing literature. To the best of our knowledge, recent empirical studies of WiMAX performance involved testbeds using only point-to-point links between a base station and a single subscriber station. In particular, we evaluated VoIP and audio/video streaming performance over a fixed WiMAX testbed in Oulu, Finland, comprising one base station and two subscriber stations.

We employed multiple competing traffic sources and measured the capacity of our WiMAX equipment to handle tens of VoIP flows between subscriber stations while delivering a variable number of video streams. We measured throughput, packet loss, and one-way delay for both line-of-sight (LOS) and non-line-of-sight (NLOS) conditions. We found out that, using best effort scheduling, the testbed can sustain 50 emulated bidirectional Speex-encoded VoIP "calls" and 5 simultaneous emulated IPTV streams, within the same WiMAX cell under LOS conditions, with negligible packet loss, adequate application-level throughput, and one-way delays within proper bounds. Under NLOS conditions, the testbed can sustain 50 emulated bidirectional VoIP flows, but only one IPTV stream with the same quality of service. We also evaluated an open source implementation of the IEEE 1588 Precision Time Protocol and compared it with GPS-assisted one-way delay measurements. We expect our results to be of great value to researchers and practitioners alike.

Acknowledgement

This work was conducted at the VTT Converging Networks Laboratory within the framework of the IST Sixth Framework Programme Integrated Project WEIRD (IST-034622), which is partially funded by the Commission of the European Union. Study sponsors had no role in study design, data collection and analysis, interpretation, or writing the report. The views expressed do not necessarily represent the views of VTT, the WEIRD project, or the Commission of the European Union. We thank the anonymous reviewers for their comments and suggestions; they greatly improved the quality of this paper.

References

[1] J. G. Andrews, A. Ghosh, and R. Muhamed, *Fundamentals of WiMAX: Understanding Broadband Wireless Networking*. Prentice Hall, 2007.

[2] IEEE 802.16 Working Group, *IEEE Standard for Local and Metropolitan Area Networks. Part 16: Air Interface for Fixed Broadband Wireless Access Systems*. IEEE Std. 802.16-2004, October 2004.

[3] ——, *IEEE Standard for Local and Metropolitan Area Networks. Part 16: Air Interface for Fixed Broadband Wireless Access Systems. Amendment 2: Physical and Medium Access Control Layer for Combined Fixed and Mobile Operation in Licensed Bands*. IEEE Std. 802.16e-2005, December 2005.

[4] K. Correll, N. Barendt, and M. Branicky, "Design Considerations for Software Only Implementations of the IEEE 1588 Precision Time Protocol," in *Proc. Conference on IEEE 1588*, Winterthur, Switzerland, October 2005.

[5] K. Correll, *PTP daemon (PTPd)*. Available: http://ptpd.sourceforge.net.

[6] International Electrotechnical Commission, *Precision clock synchronization protocol for networked measurement and control systems*. Reference numbers: IEC 61588:2004(E), IEEE 1588-2002(E), September 2004.

[7] N. Scalabrino, F. D. Pelegrini, I. Chlamtac, A. Ghittino, and S. Pera, "Performance evaluation of a WiMAX testbed under VoIP traffic," in *Proc. First ACM International Workshop on Wireless Network Testbeds, Experimental evaluation and CHaracterization (WiNTECH)*, September 2006, pp. 97–98.

[8] N. Scalabrino, F. D. Pellegrini, R. Riggio, A. Maestrini, C. Costa, and I. Chlamtac, "Measuring the quality of VoIP traffic on a WiMAX testbed," in *Proc. Third International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TRIDENTCOM)*, May 2007, pp. 1–10.

[9] O. Grondalen, P. Gronsund, T. Breivik, and P. Engelstad, "Fixed WiMAX field trial measurements and analyses," in *Proc. 16th IST Mobile and Wireless Communications Summit*, July 2007, pp. 1–5.

[10] K. Pentikousis, E. Piri, J. Pinola, F. Fitzek, T. Nissilä, and I. Harjula, "Empirical evaluation of VoIP aggregation over a fixed WiMAX testbed," in *Proc. Fourth International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TRIDENTCOM)*, Innsbruck, Austria, March 2008.

[11] S. Sengupta, M. Chatterjee, S. Ganguly, and R. Izmailov, "Improving R-Score of VoIP streams over WiMAX," in *Proc. IEEE ICC 2006*, vol. 2, June 2006, pp. 866–871.

[12] M. Chatterjee and S. Sengupta, *WiMAX Applications*. CRC Press, 2008, ch. VoIP over WiMAX.

[13] S. Mignanti, G. Tamea, I. Marchetti, M. Castellano, A. Cimmino, F. Andreotti, M. Spada, P. M. Neves, G. Landi, P. Simoes, and K. Pentikousis, "WEIRD testbeds with fixed and mobile WiMAX technology for user applications, telemedicine and monitoring of impervious areas," in *Proc. Fourth International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TRIDENTCOM)*, Innsbruck, Austria, March 2008.

[14] W. Wang, S. C. Liew, and V. Li, "Solutions to performance problems in VoIP over a 802.11 wireless LAN," *IEEE Transactions on Vehicular Technology*, vol. 54, no. 1, pp. 366–384, Jan. 2005.

[15] S. Garg and M. Kappes, "An experimental study of throughput for UDP and VoIP traffic in IEEE 802.11b networks," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, vol. 3, March 2003, pp. 1748–1753.

[16] A. Vasan and A. U. Shankar, "An empirical characterization of instantaneous throughput in 802.11b WLANs," University of Maryland, College Park, Tech. Rep., October 2002.

[17] X. Dong, M. Ergen, P. Varaiya, and A. Puri, "Improving the aggregate throughput of access points in IEEE 802.11 wireless LANs," in *Proc. 28th Annual IEEE International Conference on Local Computer Networks (LCN)*, pp. 682–690.

[18] I. Marsh, J. C. M. Severiano, V. Y. D. Nunes, and G. Maguire Jr., "IEEE 802.11b voice quality assessment using cross-layer information," in *Proc. First Workshop on multiMedia Applications over Wireless Networks (MediaWiN)*, April 2006.

[19] International Telecommunication Union, *Advanced video coding for generic audiovisual services*. ITU-T Recommendation H.264, 2005.

[20] ISO/IEC, *Information technology - Coding of moving pictures and associated audio for digital storage media at up to about 1.5 Mbit/s – Part 3: Audio*. ISO/IEC 11172-3, 1993.

[21] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, *RTP: A Transport Protocol for Real-Time Applications*. IETF RFC 3550, July 2003.

[22] S. Wenger, M. Hannuksela, T. Stockhammer, M. Westerlund, and D. Singer, *RTP Payload Format for H.264 Video*. IETF RFC 3984, February 2005.

[23] H. Schulzrinne, A. Rao, and R. Lanphier, *Real Time Streaming Protocol (RTSP)*. IETF RFC 2326, April 1998.

[24] J. Manner, *Jugi's Traffic Generator (JTG)*, Available: http://hoslab.cs.helsinki.fi/savane/projects/jtg.

[25] J.-M. Valin, *Speex: A Free Codec for Free Speech*, Available: http://www.speex.org.

[26] J. Prokkola, M. Hanski, M. Jurvansuu, and M. Immonen, "Measuring WCDMA and HSDPA delay characteristics with QoSMeT," in *Proc. IEEE International Conference on Communications (ICC)*, Glasgow, Scotland, June 2007, pp. 492–498.

[27] D. L. Mills, *Network Time Protocol (Version 3)*. IETF RFC 1305, March 1992.

[28] L. Benetazzo, C. Narduzzi, and M. Stellini, "Analysis of clock tracking performances for a software-only IEEE 1588 implementation," in *Proc. IEEE Instrumentation and Measurement Technology Conference*, May 2007, pp. 1–6.

[29] B. Wang, K. I. Pedersen, T. E. Kolding, and P. E. Morgensen, "Performance of VoIP on HSDPA," in *Proc. IEEE 61st Vehicular Technology Conference (VTC)*, vol. 4, June 2005, pp. 2335–2339.

[30] C. Bormann, C. Burmeister, M. Degermark, H. Fukushima, H. Hannu, L.-E. Jonsson, R. Hakenberg, T. Koren, K. Le, Z. Liu, A. Martensson, A. Miyazaki, K. Svanbro, T. Wiebke, T. Yoshimura, and H. Zheng, *RObust Header Compression (ROHC): Framework and four profiles: RTP, UDP, ESP, and uncompressed*. IETF RFC 3095, July 2001.