

Experimental Evaluation and Characterization of the *Magnets* Wireless Backbone

Roger P. Karrer
Deutsche Telekom Laboratories
10587 Berlin (Germany)
roger.karrer@telekom.de

István Matyasovszki
University of Limerick
Limerick (Ireland)

Alessio Botta, Antonio Pescapé
University of Napoli "Federico II"
80125 Napoli (Italy)
{a.botta,pescapè}@unina.it

ABSTRACT

High-speed wireless backbones have the potential to replace or complement wired connections. This paper provides a comprehensive network and transport layer performance evaluation of the *Magnets* WiFi backbone. The backbone, deployed in a metropolitan area of Berlin, consists of six 108 Mbps capable links using directional antennas and spans over 2.3 km. Built with off-the-shelf hardware, it features mixed 802.11a/g technology, link distances between 330 m and 930 m and support for two enhanced MAC/PHY layer modes at the access points (AP) to improve their performance. These unique characteristics provide a challenging environment to investigate the impact of a wide range of parameters. In particular, using *CBR* and *VBR* traffic profiles, we assess the influence of distance, 802.11 technology and AP modes on throughput, delay, packet loss and jitter between pairs of adjacent nodes. For example, our measurements show that the average UDP throughput of the 802.11g links varies between 15.9 and 18.7 Mbps, whereas the 802.11a link achieves 27.8 Mbps. Finally, the average UDP throughput can even be increased to 55.2 Mbps by concomitantly enabling the two enhanced AP modes.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design: Wireless communication

General Terms

Performance, Measurement, Design, Reliability

Keywords

wireless metro area networks, mesh networks, directional antenna

1. INTRODUCTION

Wireless broadband technology has recently extended its availability into the realm of access and metro area networks. It promises to combine the reliability, robustness, and wide coverage of cellular networks with the high bandwidth known from wired connections,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WiNTECH'06, September 29, 2006, Los Angeles, California, USA.
Copyright 2006 ACM 1-59593-538-X/06/0009 ...\$5.00.

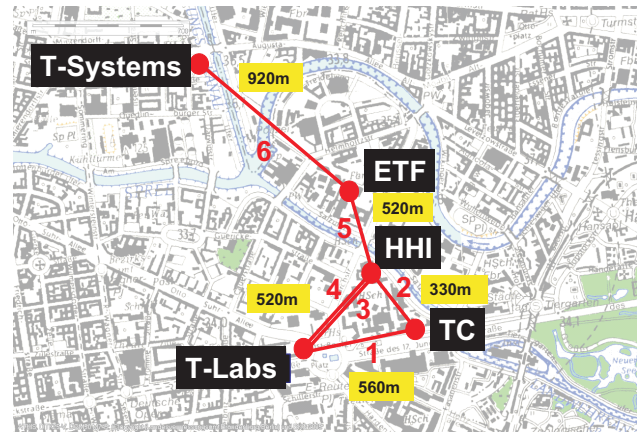


Figure 1: *Magnets* WiFi backbone in the heart of Berlin.

at a fraction of the fiber costs. Therefore, wireless broadband networks have the potential to bring us one step closer towards the vision of ubiquitous high-speed Internet access [5].

Unfortunately, our understanding of the fundamental parameters of wireless access networks, ranging from capacity constraints, scalable deployment and efficient management, up to revenue potential for Telecom Operators is still in its infancy. The primary source of knowledge about wireless networks has so far come from wireless ad-hoc / LAN networks or simulations. However, this knowledge from small-scale and slow networks does not automatically translate into metro area high-speed and heterogeneous wireless access networks where a coverage capacity of 100 Mbps or even Gbps is required over entire urban areas. Moreover, while most major cities already provide or plan to provide wireless coverage via hot spots, we often ignore how far hot-spot connectivity can be extended to a satisfactory degree. It is therefore imperative that wireless metropolitan area networks are deployed in real settings and that their fundamental behavior is assessed. Moreover, wireless access networks must be integrated into the wired Internet infrastructure to study the behavior of a wired-cum-wireless system.

This paper presents a measurement study of the *Magnets* backbone. The *Magnets* project¹ aims at deploying, analyzing and evaluating metro area heterogeneous wireless network infrastructures in the Berlin urban area. As one of the fundamental parts of the *Magnets* network, the high-speed wireless backbone connects 5 buildings in the heart of Berlin via directional antennas, spanning over 2.3 km in length, as depicted in Figure 1.

¹www.deutsche-telekom-laboratories.de/~networks/magnets.html

The objective of this measurement study is to assess the link characteristics of a wireless access network that is deployed in a dense urban area. While the APs and directional antennas *should* enable high transmission rates of 54 Mbit/s or even 108 Mbit/s, the effective throughput is a complex function of technical and environmental factors, such as interference or rain. Towards this objective, we face a rich and challenging parameter space. The link characteristics vary as a function of the distance, the Industrial Scientific and Medical (ISM) frequency band used, channel used within a band - and two configurable enhanced modes termed (*Turbo* and *Burst Mode*). Since this study focuses on a high-level link characterization, we measure the link characteristics using transport-layer tools and assess the end-to-end performance variations as a function of the transport-layer protocol (TCP, UDP), the packet size, the packet rate, the generated traffic patterns and the source-destination pairs.

For the evaluation, we first develop a measurement strategy and perform a rich set of measurements. Second, we analyze the measurements to characterize the *Magnets* backbone. In particular, for a selected parameter space, we measure and analyze round-trip time, packet loss, jitter and throughput. Our measurements show, e.g., that the average UDP throughput of the 802.11g links varies between 15.9 and 18.7 Mbps, whereas the 802.11a link achieves 27.8 Mbps. By enabling the two enhanced AP modes, the average UDP throughput can even be increased to 55.2 Mbps. Finally, the use of directional antennas, a mix of different frequency bands and separate non-overlapping channels, allows for spatial reuse and reduced interference among the backbone links.

This paper is organized as follows. Section 2 describes the wireless backbone architecture and its parameter space. Section 3 describes the measurement methodology. Section 4 presents and discusses the obtained results. After discussing related works in Section 5, we conclude the paper in Section 6.

2. MAGNETS WIFI BACKBONE

2.1 Backbone Setup

The *Magnets* architecture consists of three basic network parts: (i) a high-speed 802.11-based *wireless backbone* consisting of point to point links; (ii) an 802.11-based wireless access *mesh network*; and (iii) a *heterogeneous access network* that combines multiple wireless technologies, such as GPRS, UMTS, WiFi and WiMax [6].

The objective of the WiFi backbone is to assess the feasibility and limitations of replacing a wired networking infrastructure by a high-speed wireless backbone. Such a replacement can be envisioned for two main reasons. First, wireless connectivity is suitable in areas where fiber is currently not available (in fact, the eastern part of Berlin still does not provide DSL connectivity to all homes), in rural areas or in developing countries. Second, virtual providers that offer connectivity but do not own an entire network currently lease lines from operators. A wireless backbone may replace the need to lease lines.

The details of the *Magnets* WiFi backbone deployment in Berlin are depicted in Figure 2. All directional antennas reside on top of high-rise buildings and have unobstructed line of sight. The distances between the buildings varies between 330 m and 930 m, resulting in a total span of approximately 2.3 km between T-Labs and T-Systems. All transmissions are in the unlicensed ISM spectrum (2.4 and 5 GHz) and all backbone components (antennas, access points) consist of off-the-shelf hardware supporting 802.11a/g modes at 54 Mbps. A proprietary protocol enhancement called *Turbo Mode* can be optionally switched on to double the gross throughput to 108 Mbps. Therefore, the *Magnets* backbone is a first step towards a realization of a 100x100 type access network [15].

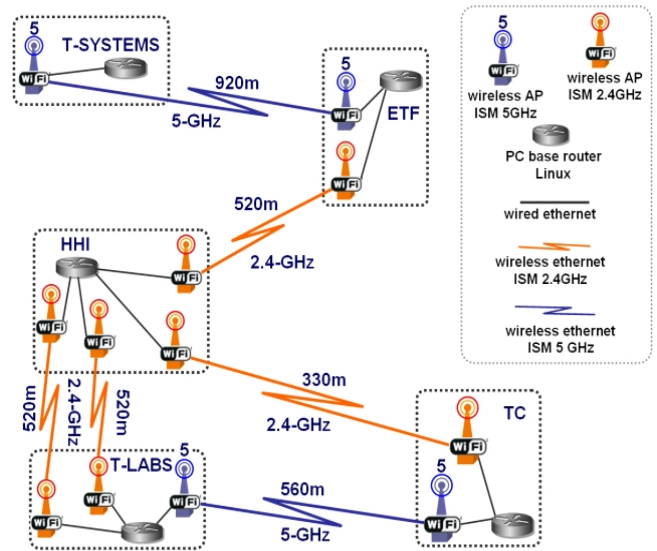


Figure 2: *Magnets* WiFi backbone.

Each *Magnets* backbone node consists of an IntelP4-PC based router with a 3 GHz processor, 1 GB of RAM and 80 GB HDD running Linux with kernel version 2.6.15. We opted for a PC-based solution to be able to install computationally intensive network management and application-specific services on these nodes. Backbone routing is done using OSPF². Each backbone router has attached a number of LanCom WiFi access points (APs) and on total there are 12 APs in the backbone. 10 of the 12 APs are mounted outdoors (OAP54) to minimize the cable length to the antennas⁴, while at the ETF node APs (IAP54) were placed indoors and antennas outside. To minimize antenna cable length at ETF, both indoor APs and their corresponding outdoor antennas were symmetrically positioned and wall-mounted.

Of particular interest is that 2 parallel links are deployed between T-Labs and HHI. These parallel links (referred to as “twin links”) are independent as they have their own APs and antennas and the distance between their respective antenna centers on the masts is roughly 50 cm. Routing can be configured to consider the twin links as two distinct paths or to load-balance traffic flowing through them.

The access points are based on Intel IXP420@266 MHz (IAP) and IXP425@533 MHz (OAP) programmable network processors (NPs) respectively. APs based on the IXP4xx series architecture have been shown to outperform APs based on Broadcom or IBM processors from the same class [13]. The APs run LC.OS [8], a proprietary OS based on either VxWorks or Linux and are powered via Power over Ethernet (PoE). AP WLAN interfaces are based on the Atheros 5213/5112 chipset⁵. The access points are connected to directional antennas. 8 antennas operate at 2.4 GHz with a gain of 18 dBi, 4 operate at 5 GHz with 9°/23 dBi. The antennas along with the APs are mounted on special steel masts that hold up to 4 antennas and APs and resist the wind load of the attached devices.

²<http://www.quagga.net>

³<http://www.lancom-systems.de>

⁴Deployment pictures at <http://www.deutsche-telekom-laboratories.de/karrer/magnets.html>

⁵<http://www.atheros.com>

Table 1: Magnets Backbone and Traffic Parameter space.

Level	Parameter	Values
Link	Frequency	2.4 and 5 GHz
	Channel	3 and 19 orthogonal channels
	<i>Turbo Mode</i> <i>Burst mode</i>	on/off on/off
Topology	src-dst	any of the 5 nodes
	interference	single link / all links
	hop length	1 – 6 hops
Traffic	Pattern	CBR, VBR
	Rate	$\in \{1Mbps, 90Mbps\}$
	Packet size	1024 Bytes
	Protocol	TCP, UDP

2.2 Backbone Parameters Space

The *Magnets* backbone provides a large parameter space that can be separated into 3 groups, as shown in Table 1: link parameters, topology parameters and traffic parameters.

The first group refers to parameters that can be set on a per-link basis. The *Magnets* APs support both 802.11a and 802.11g technology, thus frequencies in the 2.4 and 5 GHz can be used. However, the antennas support only either the 802.11a or the 802.11g frequency band, so that a change in the spectrum is only possible if a different antenna is mounted. Within the 2.4 GHz band, 3 orthogonal channels exist, whereas 19 are available in the 5 GHz range. We performed initial measurements to determine the configuration with the least interference. The channel assignment is not modified during our experiments. Finally, the APs have two low-layer modes that can be enabled optionally: a *Turbo Mode* and a *Burst Mode*[7]. The *Turbo Mode* doubles the transmission rate by enlarging the channel from 20 MHz to 40 MHz. In the 2.4 GHz, the *Turbo Mode* frequency is centered around channel 6, using a spectrum between 2417 MHz to 2457 MHz. Due to its innate features, *Turbo Mode* interferes with all channels in the 2.4 GHz range. However interference is alleviated due to the use of directional antennas. In the 5 GHz range, 3 orthogonal Turbo Mode channels are available in the lower band range and 2 channels in the upper band range. The *Burst Mode* enables an AP to increase its sending rate. In contrast to “normal” mode where the sender has to wait a certain period equal to the Distributed Inter-Frame Space (DIFS) after receiving an ack until it can send a new packet. In Burst Mode, the sender only waits for the shorter SIFS (Short Inter-Frame Space) period. While *Burst Mode* leads to severe unfairness in general wireless environments, it enables *Magnets* to increase the throughput on the individual links significantly without negative impact on neighboring links due to the spatial reuse and the directional antennas.

As for the second group, the choice of traffic source and destination allow a variation of the network topology. Depending on the sender and receiver, the number of hops can be varied from 1 to 4. The path length can even be increased to 6 if a loop is introduced between T-Labs, HHI, TC and back to T-Labs and the routing on the parallel links between T-Labs and HHI are configured accordingly. Finally, measurements can be performed by activating only a single link or by measuring multiple links in parallel. In the latter case, the links may interfere.

Finally, the measurement traffic injected into the network can be varied in its pattern, the sending rate and the packet size. We assess the backbone performance for UDP and TCP. More details on the parameter values chosen for the evaluation are given in the following section.

3. METHODOLOGY AND MEASUREMENT APPROACH

We use an active measurement approach to investigate the performance of the *Magnets* backbone evaluating the statistics of four QoS parameters (throughput, jitter, delay, and packet loss). A software platform called D-ITG (Distributed Internet Traffic Generator) [11] generates synthetic traffic between two chosen end points. D-ITG is composed of five blocks: sender (ITGSend), receiver (ITGRecv), log server (ITGLog), manager (ITGManager), and decoder (ITGDec, used to analyze the produced log files in order to extract the QoS parameter statistics). D-ITG is capable of producing a large variety of random traffic profiles by modeling the *Inter Departure Time* (IDT) and *Packet Size* (PS) processes. Moreover, it is able to log the sent and received packets at, respectively, sender and receiver sides. This allows to establish a relationship between an experimentally observed performance value and its cause. For TCP traffic, we extract the throughput, round-trip time and jitter, for UDP we additionally measure the packet loss.

We report round-trip time (RTT) rather than one-way delay because the available clocks on the routers and APs are not accurately synchronized to derive exact measurements. We measure the TCP and UDP RTT in the same way. More precisely, with the term RTT we mean the following: for each packet we measure the time between sender/receiver plus the time between receiver/sender.

For the measurements we restrict the topology to 3 backbone nodes: T-Labs, HHI and TC. This topology provides several interesting aspects for the evaluation. First, the twin links allow a characterization of the influence of environmental factors. Second, the links TLabs - HHI and TLabs - TC both span a distance of 500+ meters, but differ in technology (802.11g / 802.11a respectively). Finally, performance improvements achieved by simultaneously enabling *Turbo* and *Burst Mode* enhancements are analyzed.

We have performed a large set of experiments with different traffic profiles. The profiles differ from each other in terms of bitrate as well as PS and IDT patterns. Being the least random, the Constant Bit Rate (CBR) traffic can be useful to draw a reference line for the statistics collected using Variable Bit Rate (VBR) traffic. For this reason and to limit the number of variable parameters that must be taken into account, we present the results related to two CBR traffic profiles denominated as *low* and *high* traffic. The *low* traffic profile consists of 1024 bytes packets sent at a frequency of 128 pkt/s, corresponding to an offered bitrate of 1048 Kbps. For the *high* traffic profile, 1024 bytes packets are injected at a rate of 11000 pkt/s, corresponding to an offered bitrate of 90 Mbps. Details related to the traffic profiles are shown in the third row of Table 1. While the *low* traffic profile is far from the saturation point, we found through initial measurements that the backbone nodes are able to generate up to 11000 pkt/s. At this rate, we are above the wireless channel saturation point. The *high* traffic profile characterizes the backbone performance under saturation.

For the selected set of parameter values, traffic was generated for a duration of 120 seconds. This duration has been chosen to correctly evaluate TCP performance out of its transient phase. All packets sent and received are timestamped and logged. The round trip time, jitter and packet loss samples are directly calculated from logs (on a per packet basis) and then averaged over intervals of 50 milliseconds. For bandwidth, the amount of bytes transmitted or received in 50 milliseconds interval is summed up. In order to minimize the effect of random errors, we performed 20 (interleaved) test repetitions and averaged the results. During the measurement stage, between May and June 2006, 40 GB of data traces were collected and they are freely available at [10].

4. EXPERIMENTAL ANALYSIS

In this section we report the most interesting results obtained from our measurements. A complete overview of results is reported in [3].

4.1 UDP Results

First, we assess the throughput of the different links measured with UDP. Since the links have a raw capacity of 54 Mbps, it can be expected that the transport-layer goodput reaches up to 30 Mbps. The corresponding results are shown in Figure 3.

Figure 3(a) shows the throughput under *high* traffic conditions using UDP. The figure first shows that the link TLabs-TC has the highest throughput and the lowest variation: the average bandwidth reaches 27.8 Mbps and the standard deviation is 2.14 Mbps. Since the throughput is higher than that of links with similar distance (TLabs-HHI) and longer than HHI-TC, we attribute the high throughput to the frequency choice. In particular, the Dynamic Power Selection (DPS) of 802.11a dynamically adjusts the sending power. The sending power can thereby exceed that of 802.11g links and the dynamic adaptation leads to stable throughput conditions. The DPS thereby compensates the higher damping of the signal at 5 GHz compared to 2.4 GHz.

The throughput of the other links varies between 15.9 Mbps and 18.7 Mbps. The difference among the links is not significant. Nevertheless, the measurements show some interesting trends. First, the short link TC-HHI has a lower average throughput and a higher standard deviation than the links between TLabs and HHI. There are two potential explanations. Even though both links have perfect line of sight, TC has a lower height (9 stories) than TLabs (20). Therefore, interference from neighboring APs may have a slight, yet measurable interference. Second, link TC-HHI is assigned channel 6, whereas TLabs-HHI are set to 6 and 13 respectively. The interference may vary as a function of the assigned channel.

Figure 3(b) shows the jitter distribution of the different links. The 5 GHz link has a lower jitter (less than 8 milliseconds) with small variance, whereas HHI-TC and TLabs-HHI-2 have higher values and a broader distribution. The twin links again show a non-consistent behavior in spite of their parallel deployment and the fact that they use the same hardware. Thus, our conclusions here are three-fold. First, the use of DPS at 5 GHz has a significant impact on jitter as well. Second, the distance has a negligible impact, as the jitter is comparable between HHI-TC and TLabs-HHI-1 in spite of their difference in distance (330m vs. 520m). Third, in spite of a parallel deployment, the jitter of the twin links is significantly different. We attribute the difference again to environmental factors, but a more detailed investigation is needed to confirm this hypothesis and to quantify the impact of these factors.

Figure 3(c) depicts the round-trip time distributions. Here, we note a clear difference among the RTT as a function of the link distance. In particular, link HHI-TC which spans 330m has an average of 21 ms with a standard variation of $8.26e - 003$ s. In contrast, the links that span 520 and 560m have average values up to 51.7 ms. More precisely, the links TLabs-TC and TLabs-HHI-1 have a similar distribution (with average values of 29.2 and 32.9 ms respectively) whereas TLabs-HHI-2 has a clearly higher average of 51.7 ms.

Finally, Figure 3(d) depicts the packet loss distribution. The link at 5 GHz slightly outperforms the links at 2.4 GHz: TLabs-TC has an average of 7600 pkt/s, whereas the 2.4 GHz links have an average packet loss of 8900 pkt/s. To understand the number, remember that traffic is injected at 11000 pkt/s with size of 1024 bytes, which corresponds to an offered load of 90 Mbps.

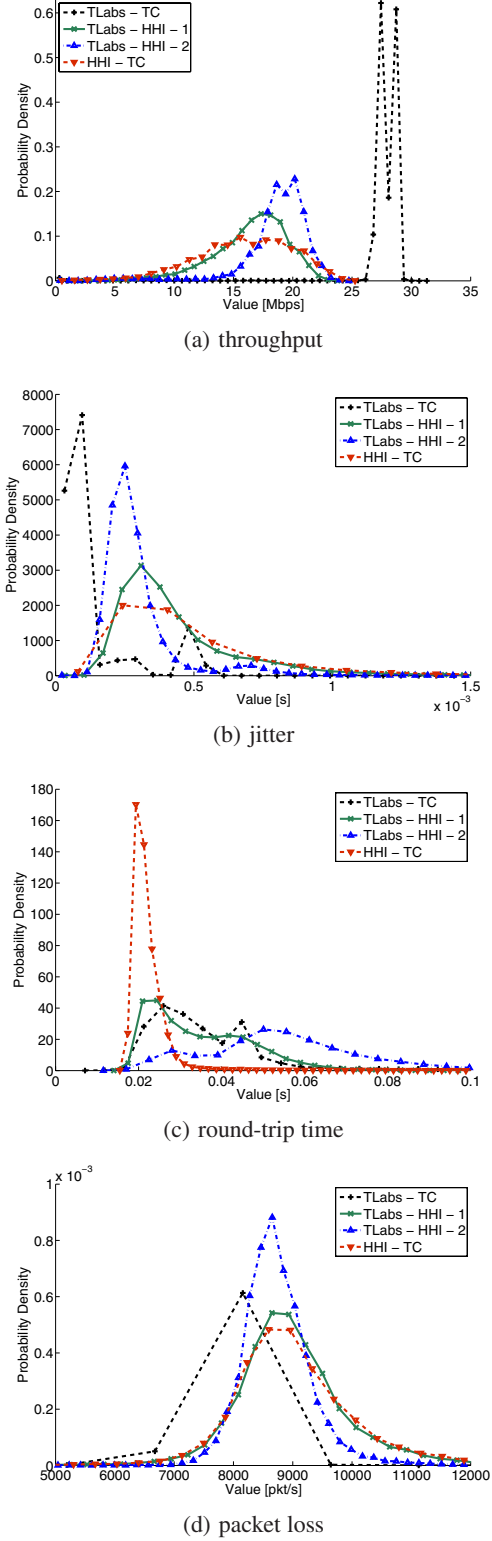


Figure 3: PDFs of *high* UDP traffic.

With a sustained throughput of 25 Mbps, only 35% of the packets can be transmitted successfully. Therefore, the measured packet loss corresponds to the observed throughput. In addition to the

measurements carried out at *high* traffic conditions, we additionally performed measurements under *low* traffic conditions. Given an injected traffic rate at 1 Mbps, we also measured an average throughput of 1 Mbps with insignificant derivations. Also the jitter values are below the threshold to draw clear conclusions. Therefore, the corresponding figures are not shown here.

4.2 TCP Results

Here, we measure the throughput using TCP. These measurements enhance the UDP measurements as the measured throughput is additionally influenced by the end-to-end congestion control of TCP. The corresponding results are shown in Figure 4.

Figure 4(a) shows the throughput of the different links under *high* traffic condition using TCP. First and foremost, we note that the throughput characteristics is similar to that under UDP conditions in Figure 3(a). The 5 GHz link (TLabs-TC) clearly stands out and reaches an average throughput of 25.8 Mbps and a standard deviation of 1.56 Mbps. Thus, compared to UDP, the throughput differs only by 2 Mbps.

The links at 2.4 GHz have an average of 12 – 16 Mbps and their standard deviation varies between 2.71 and 3.45 Mbps. These results are as expected given the UDP results. In particular, consider TCP models that model TCP throughput as a function of the RTT and packet loss[9]. For the 5 GHz link between TLabs and TC, the lower packet loss ratio leads to a higher TCP throughput of 25.8 Mbps. Moreover, the throughput is relatively stable with a standard deviation of 1.56 Mbps. Next, comparing the 2.4 GHz links, we note a higher average throughput for the link between HHI and TC. Since all links have a similar packet loss distribution in the UDP measurements, this effect is explained with the lower RTT of the shorter link (see also Figure 4(c)). Therefore, the average throughput of the link HHI-TC is 16.6 Mbps and the standard deviation 3.08 Mbps. Finally, note that the throughput of the twin links is now almost equal. The measured differences in the RTT have a negligible impact on the TCP throughput.

Figure 4(b) depicts the PDF of the jitter. We measured the jitter because more and more multimedia applications are deployed that use TCP as their underlying protocol. First, note that all jitter values are in the order of 1 ms, making a very good case for the use of multimedia applications. Among the links, the 5 GHz link again outperforms the 2.4 GHz links. Among the latter, little difference is seen between TLabs-HHI-1 and HHI-TC, leading to the conclusion that the link distance has a negligible impact on jitter. In contrast, the twin links show again a slightly different behavior but it is worth noting that they have very similar average values ($1.09e - 003$ and $1.29e - 003$ s).

Finally, Figure 4(c) depicts the round-trip time. The short link between HHI and TC has a significantly shorter RTT than the links that span 500+m, but its distribution is bimodal with one mode closer to 0. At such small time scales, OS influences such as buffering at the kernel can be the cause. In contrast, for the other links, the distribution is gaussian with average values ranging from 0.205 to 0.349 s.

Besides the measurements under *high* traffic conditions, we performed measurements under *low* traffic conditions. As for the UDP measurements, the throughput results show little variations and are close to 1 Mbps. The RTT and jitter measurements are equally identical for all links. Therefore, these measurements can be used as a calibration for the testbed and confirm the validity of our results. However, a quantitative analysis is difficult, as the measured values are very low. Therefore, it is not possible to clearly identify whether the observed values are caused by the wireless links or by other effects, such as buffering in the OS.

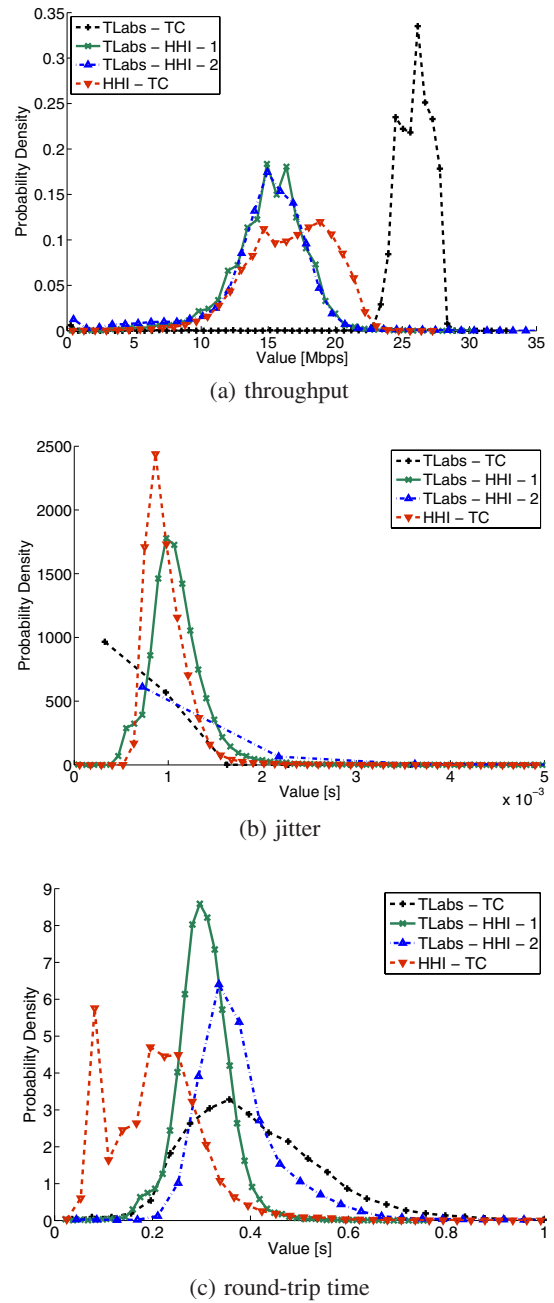


Figure 4: PDFs of *high* TCP traffic.

In conclusion, the measured values and the observations made are consistent. The UDP measurements can be explained, even though further research is needed to confirm and to quantify the impact of environmental factors. Moreover, the results of the TCP measurements can be derived from the UDP results. Therefore, we argue that we have collected a consistent set of measurements and that the presented characterization is valid.

The stability of the 5 GHz has come as a surprise as the damping at 5 GHz should have had a quantitative impact on throughput. However, the Dynamic Power Selection has more than compensated the physical drawback and lead to a link with higher throughput and lower throughput variation. These characteristics can have

significant advantages for applications in a real deployment. First, for applications such as adaptive multimedia applications, the high average throughput and the low variations lead a better streaming quality. Second, if traffic engineering is required along the backbone, such as providing QoS for VoIP or TV, or to shape flows to provide fairness and spatial reuse, the stability of this link may significantly simplify algorithms and lead to better performance.

4.3 Temporal Structures

Here, we assess the temporal structures of the collected time series - via the autocorrelation function (ACF) - with the aim to compare its behavior in *low* and *high* traffic conditions. The analysis of the ACF provides hints regarding time series proprieties. If high values or periodical trends are evidenced, prediction could be used to estimate the parameter behavior.

Figure 5 shows the autocorrelation function (ACF) of the throughput and jitter in *low* traffic conditions with UDP and TCP. In the case of TCP (Figures 5(a) and 5(b)), the values oscillate around their mean value. The same results for UDP throughput (Figure 5(c)). Moreover, the behavior of the different links is similar throughout all figures with the periodic spikes attaining high autocorrelation values. This reveals that the time series repeat itself almost unchanged. Therefore, using prediction techniques, future values could be estimated. In contrast, the UDP jitter (Figure 5(d)) reveals a hybrid behavior: the values oscillate, though at a lower amplitude, and the mean shifts towards zero. Therefore, the correlation among the samples is low.

Figure 6 depicts the ACFs of the throughput and jitter under *high* traffic conditions. In contrast to the *low* traffic conditions, the throughput characteristics significantly differ among the links. First, the 5 GHz link from TC to TLabs has an oscillating pattern similar to the one under *low* traffic conditions. In contrast, the links in the 2.4 GHz spectrum have a decreasing autocorrelation. Finally, there is considerable variation among the 2.4 GHz links. In Figure 6(a), e.g., the short link between HHI and TC has a low-level frequency oscillation whereas the twin links have an almost linear low-level component overlapped with a high-frequency component.


Finally, the TCP jitter autocorrelation follows the TCP throughput, whereas the UDP jitter shows very little autocorrelation. Interesting is the high autocorrelation value of the 5 GHz link that reveals a very stable jitter.

4.4 Performance with VBR Traffic


Internet traffic is composed of a mix of heterogeneous traffic types. Besides the *low* and *high* CBR traffic, we tested the performance of the link HHI-TC with VBR traffic. Following the video traffic model proposed in [14], a stream is sent at a rate equal to 720 pkt/s (with a constant profile) and packet sizes following a normal distribution with a mean of 926, 4 bytes and standard deviation of 289, 5 bytes. The rate of 720 pkt/s corresponds to a frame rate of 24 frames per second, and, assuming 30 slices per frame, we have a single slice per packet.

Figure 7(a) shows that the throughput measured on the link follows the imposed rate of the sender (video source). The packet loss probability (not shown) is basically equal to zero. Moreover, the jitter measured in this case (shown in Figure 7(b)) is lower than that one obtained over the same link in the *low* traffic load condition.


These results confirm that - in absence of cross-traffic - the observed link is able to carry a video transmission without performance degradation. This is encouraging, as previous measurements of video streaming over wireless ad-hoc networks raised the question whether wireless technology could support video streaming.




(a) TCP throughput



(b) TCP jitter




(c) UDP throughput




(d) UDP jitter


Figure 5: ACFs of *low* traffic.




(a) TCP throughput



(b) TCP jitter



(c) UDP throughput



(d) UDP jitter

Figure 6: ACFs of high traffic.

4.5 Performance with Turbo and Burst Mode

One the main contribution of this work is the analysis of the performance gain obtained with two enhanced MAC/PHY layer modes at the access points (*Turbo Mode* and *Burst Mode*). Figure 8 compares the throughput and round-trip time of the link TLabs-HHI-1 with and without *Turbo Mode* and *Burst Mode*. First, Figures 8(a) and 8(b) show their impact on TCP under *high* traffic conditions. The average throughput increases from 15.1 Mbps to 48.5 Mbps,




Figure 7: VBR Video streaming over the backbone.

a 3.2-fold increase. While the *Turbo Mode* doubles the throughput by doubling the channel width, the *Burst Mode* enables higher sending rates by decreasing the wait time after a transmission. This decrease also leads to a shorter round-trip time, as shown in Figure 8(b). In particular, the average round-trip time decreases from 243 ms (as in Figure 4(c)) to 130 ms - a net reduction of 53%. Moreover, the variation in the round-trip time is smaller, resulting in more stable end-to-end performance.

Figures 8(c) and 8(d) show the corresponding results for UDP. The average throughput increases even to 55.2 Mbps, a net gain of 38.7 Mbps. The average round trip time is reduced from 33 ms without *Burst Mode* to 20 ms with *Burst Mode*. Again, the lower value distribution leads to a more stable environment.

4.6 Discussion

The performance evaluation of the Magnets WiFi backbone shows that high-speed wireless connectivity is feasible. The deployment and the evaluation of a backbone in a metropolitan area is utterly challenging and the achieved performance is beyond our expectation given the potential influence of the environment. For example, at HHI we found 3 competing networks at channel 1, 5 and 2 at channels 6 and 11 respectively. At ETF, we even found 26 competing networks at channel 6. Additionally, the backbone link side lobes have the following radiuses: both TLabs-HHI links have a *Fresnel zone 1* radius of 4m, TLabs-TC 2.64m, TC-HHI 3.2m, HHI-ETF 4m and ETF-TSsystems 3.4m. The interference from competing APs as well as side lobes, must not be ignored or underestimated, but their impact is hard to quantify, especially that the number of competing networks is not a constant. We attribute some measured effects to these factors, such as wide distributions of a specific metric on a single link. To clearly identify the impact of interfering networks, day and night effects (reflections), temperature and humidity, a long-term performance study of the Magnets backbone will be performed.

In contrast to previous studies that painted wireless communication as unreliable, this study shows that a well-planned deployment of wireless technology is able to provide reliable connectivity. Therefore, wireless backbones are an attractive and cost-efficient alternative to wired links.

5. RELATED WORK

Performance studies of outdoor wireless backbones are extremely rare today as data from less than a handful of deployed networks is available. A large body of work has focused on indoor environments, whereas *Magnets* is designed as an access network. Moreover, the *Magnets* backbone is deployed in a dense urban area. In contrast, related backbones such as the *Digital Gangetic Plains (DGP)* [1] [12] or the *TjA* network [4] are deployed in a rural areas of India and a sparsely-populated residential area in Houston,

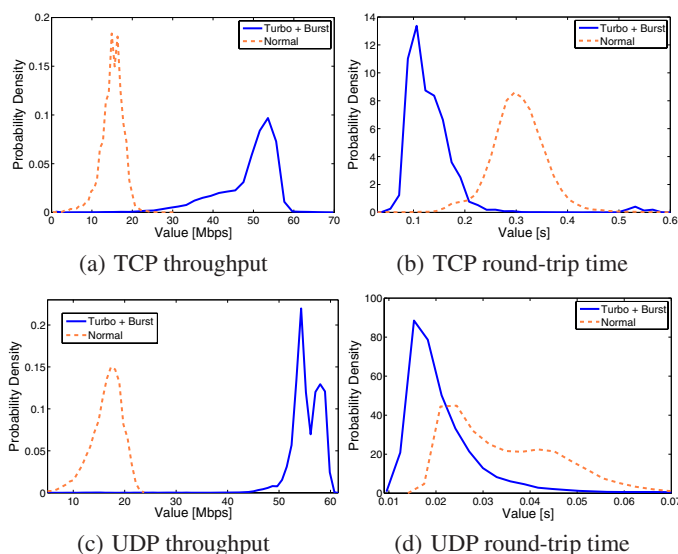


Figure 8: Impact of Turbo and Burst Modes.

Texas, respectively. The *MIT roofnet* [2] is deployed in an urban area, but only contains 3 directional antennas that are not evaluated in detail. *Magnets* provides a wide parameter space for investigation: 2.4 and 5 GHz links that span between 330 m and 920 m, with the optional enabling of *Turbo* and *Burst Mode*. Previous studies have a significantly smaller parameter space. The *MIT roofnet*, *TjA* and *DGP* operate in the 2.4 GHz only. *DGP* operates over larger distances. Also, none of these networks achieve the high rates featured by *Magnets*. Finally, we make available the data traces we collected [10] and on which this work is based.

6. CONCLUSIONS

The deployment of networks is fundamental to increase the understanding of wireless communications. The *Magnets* backbone provides a unique opportunity to study various parameters in a dense metropolitan area. This paper studies the impact of distance, frequency, topology, and traffic on the characteristics of wireless channels and its impact on higher layer protocol performance. Using off-the-shelf hardware with 54 Mbps raw throughput, we measured an average transport-layer throughput of up to 18.7 Mbps over 560 m.

In the 5GHz range where the band suffers less from interference, the average sustained throughput even reached 27.8 Mbps. Finally, by enabling low-layer enhancements, the throughput can even be improved to 55.2 Mbps. In contrast to previous studies that painted wireless communication as unreliable, this study shows that a well-planned deployment of wireless technology is able to provide reliable connectivity. Therefore, when needed, wireless backbones are an attractive and cost-efficient alternative to wired links.

In the future, we will further extend our measurements with two objectives. First, this paper covers a small part of the *Magnets* parameter space. For example, measurements of 3- to 6-hop flows, channel selection or routing impacts are not addressed yet. Second, while the paper characterizes the variations in terms of throughput, delay, packet loss, and jitter, our understanding of why these variations occur and how they can be accounted for are still in its infancy. Future measurements on the *Magnets* backbone will contribute towards understanding wireless communications in general and high-speed wireless backbones in particular.

7. ACKNOWLEDGMENTS

This work has been partially supported by PRIN 2004 Quasar Project, by CONTENT NoE and NETQOS EU project. Istvan Matyasovszki was an intern at the Deutsche Telekom Laboratories between March and May 2006.

8. REFERENCES

- [1] P. Bhagwat, B. Raman, and D. Sanghi. Turning 802.11 inside-out. In *Proceedings of HotNets-II*, Cambridge, MA, Nov. 2003.
- [2] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of the mit roofnet mesh network. In *Proceedings of ACM Mobicom'05*, Cologne, Germany, Aug. 2005.
- [3] A. Botta, I. Matyasovszki, A. Pescapé, and R. Karrer. Performance evaluation of magnets wireless backbone. Technical report, Deutsche Telekom Laboratories and University of Napoli "Federico II", July 2006. <http://www.grid.unina.it/Traffic/pub/TR-DTLab-UoN-2006.pdf>.
- [4] J. Camp, J. Robinson, C. Steger, and E. Knightly. Measurement driven deployment of a two-tier urban mesh access network. In *Proceedings of ACM MobiSys 2006*, Uppsala, Sweden, June 2006.
- [5] R. Karrer, A. Sabharwal, and E. Knightly. Enabling large-scale wireless broadband: the case for TAPs. In *Proceedings of HotNets-II*, Cambridge, MA, Nov. 2003.
- [6] R. Karrer, P. Zerfos, and N. Piratla. Magnets - a next-generation access network. In *Poster at IEEE INFOCOM'06*, Barcelona, Spain, Apr. 2006.
- [7] Lancom Systems, Germany. 108 Mbits Super A/G. <http://www.lancom-systems.de/fileadmin/produkte/features/techpaper/TP-108Mbits-EN.pdf>.
- [8] Lancom Systems, Germany. Lancom business networking. http://www.lancom-systems.de/fileadmin/produkte/feature/brochures/PO_EN_Web.pdf.
- [9] J. Padhye, V. Firoiu, D. Towsley, and J. Krusoe. Modeling TCP throughput: A simple model and its empirical validation. *Proceedings of the ACM SIGCOMM '98 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 303–314, 1998.
- [10] A. Pescapé and al. Data traces. <http://www.grid.unina.it/Traffic/Traces/Magnets.php>.
- [11] A. Pescapé and al. Distributed internet traffic generator. <http://www.grid.unina.it/software/ITG>.
- [12] B. Raman and K. Chebrolu. Design and evaluation of a new MAC protocol for long-distance 802.11 mesh networks. In *Proceedings of ACM Mobicom'05*, Cologne, Germany, Aug. 2005.
- [13] The Tolly Group. Network processors: Performance analysis of 802.11 broadband routers and access point, 2004. www.intel.com/design/network/products/npfamily/tolly_204131.pdf.
- [14] W. Willinger and M. Garrett. Analysis, modeling, and generation of self similar VBR video traffic. In *Proceedings of SIGCOMM'94*, London, England, Aug. 1994.
- [15] H. Zhang and al. 100x100 clean slate project. <http://100x100network.org/>.