

High-speed backhaul networks: Myth or reality? [☆]

Roger P. Karrer ^a, Alessio Botta ^b, Antonio Pescapé ^{b,*}

^a *Deutsche Telekom Laboratories, TU Berlin, 10587 Berlin, Germany*

^b *Dipartimento di Informatica e Sistemistica, University of Napoli "Federico II", Via Claudio, 21, 80125 Napoli, Italy*

Available online 26 January 2008

Abstract

Wireless technology promises a realization of the long-standing vision of ubiquitous high-speed Internet access. WiFi-based wireless mesh networks that provide user access and wireless data transport over a multi-hop *backhaul* network are a promising incarnation of the above vision. However, while WiFi is successfully used to provide user connectivity via access points, we note that currently deployed wireless mesh networks show a dismal performance and lack mechanisms in the backhaul to provide an efficient and fair data transport over multiple hops. To assess the capabilities and the limitations of wireless backhaul networks, we are currently building *MagNets*, a next-generation wireless mesh network in the city of Berlin. Using *MagNets*, this paper provides insight on how to plan and design efficient wireless backhaul networks by describing the work breakdown and the lessons learned from the design and deployment process. Then, we perform a comprehensive performance evaluation to investigate the impact of a wide range of parameters to shed light on the potential and limitations of wireless backhaul networks.

© 2008 Elsevier B.V. All rights reserved.

Keywords: Multi-hop wireless network; Wireless backbones; Performance measurements

1. Introduction

Wireless technology promises a realization of the long-standing vision of ubiquitous high-speed Internet access. Therefore, it has the potential to foster unforeseeable communication possibilities among humans and machines and to revolutionize society in a way the processor or the Internet did in the last century. Wireless mesh networks [8,11] are a key enabler of this vision by combining three main advantages and therefore excel over wired, cellular and other forms of wireless networks, such as ad-hoc networks:

- *Cost effectiveness:* Only a subset of the mesh nodes are required to have a fixed Internet connection. Therefore, capital and operational expenditures (CAPEX and

OPEX) are significantly reduced compared to wired and hot spot infrastructures [20].

- *Reliability and performance:* Compared to ad-hoc networks, mesh networks are primarily infrastructure-based and power wired. Therefore, they can be engineered, planned and deployed to avoid many of the limitations known from ad-hoc networks.
- *Unlicensed deployment:* In contrast to cellular networks, today's WiFi mesh networks operate in the unlicensed spectrum. They do not incur license costs and acquisition time, and can therefore be deployed by individuals and communities. As an example of a community network, the Berlin Freifunk Network currently consists of 800 mesh nodes and provides access to a large community of users in East Berlin [16].

How far are mesh networks from fulfilling the vision of a ubiquitous high-speed Internet access? The answer is not straightforward, but instead depends on the structure and the functionality of the mesh network.

[☆] This work has been partially supported by EU Content NoE.

* Corresponding author. Tel.: +39 0817683856.

E-mail addresses: roger.karrer@telekom.de (R.P. Karrer), a.botta@unina.it (A. Botta), pescap@unina.it (A. Pescapé).

Wireless mesh networks can typically be separated into two tiers: an *access* and a *backhaul* network, as depicted in Fig. 1. The access network consists of wireless mesh nodes (letters A to H) that provide connectivity to the associated users within their transmission range (circles around the mesh nodes). Users that connect to these access points may be stationary (homes) or mobile. Today's off-the-shelf hardware have a performance of 54 Mbps raw throughput by the 802.11a/g standard, which results in a net throughput of approximately 27 Mbps accounting for the overhead of the lower layers (physical and media access control (MAC)). Moreover, next-generation 802.11n equipment is projected to achieve up to 600 Mbps. Even though this capacity is shared by multiple users, the access rates exceed many of today's wired connections (DSL). Therefore, we argue that the access tier is able to fulfill the above vision.

In contrast, the backhaul network interconnects the APs (dotted lines) and transports data from and to those mesh nodes that are equipped with a wired Internet connection (e.g., DSL, node D). Mesh nodes may offer only backhaul functionality and, e.g., be mounted on lamp post (nodes F and G), or they may jointly offer access and backhaul functionality, ideally though via different wireless interfaces. The main objective of the backhaul is to forward data along the backhaul network from users to a wired mesh node and back. This objective implies that the data transmission must be efficient and fair over multiple wireless hops that typically span larger distances to limit the number of hops. In terms of performance, experiences from deployed mesh networks paint a dark picture. Assuming that most user traffic is sent between the APs and the Internet, backhaul links should at least match the capacity of the APs, but ideally be multiples of the AP capacity. However, the backhaul links in currently deployed meshes, such as the TFA network in Houston [15] or the MIT roofnet [12], show only single-digit throughputs. Thus, the effective throughput is only a fraction of the nominal capacity, and also only a fraction of the throughput created by a single access point. This dismal performance is often aggravated when the traffic is forwarded over multiple hops [17].

Where does this “digital divide” between efficient access and dismal backhaul come from? Is it an inherent problem

that wireless backhaul networks are doomed to slow rates? Do we therefore have to bury our vision of high-speed all-wireless networks? If not, how do we build high-speed wireless backhaul networks? To answer these and more questions, we are currently deploying *MagNets*, a next-generation high-speed wireless testbed in the city of Berlin. The current deployment plan foresees 100 mesh nodes with heterogeneous technology, in part connected via a high-speed wireless backbone. This paper describes its design, our experiences with the deployment and its performance evaluation via extensive experimental measurements. The main conclusion is that wireless backhaul networks can indeed achieve transmission rates of 30 Mbps over single links, and performance does not need to degrade for multi-hop traffic. These rates can even be improved with proprietary PHY and MAC layer enhancements available in modern off-the-shelf mesh nodes to double the throughput. More good news is that the backbone is hardly affected by environmental factors, such as weather or time-of-day effects. The big “however”, though, comes from interference: with up to 25 competing wireless networks at some APs on a single channel, we experienced significant performance degradations. Given the rapid deployment of WiFi networks in homes and public places, we conclude that technical solutions or political solutions (spectrum management) are urgently required towards realizing the vision of a ubiquitous high-speed wireless Internet.

This paper is organized as follows. In Section 2, we describe the architecture of the *MagNets* backbone that forms the backhaul of the *MagNets* testbed: how we planned the backhaul network, what hard- and software we used, how the locations to deploy the backbone are chosen. In addition, we discuss the potential and the complexity of an experimental evaluation of the backbone and we show that a significant number of parameters at various layers interact with each other. To understand the behavior of the backbone and to provide clear answers to the questions raised in this work, we develop a concise methodology to measure the various parameters in isolation. Then, Section 3 provides initial measurement results that show the performance results of the backbone over single and multiple hops. Moreover, besides the simple link measure-




Fig. 1. A wireless mesh network, composed of an access network (circles) and a backhaul network (dotted lines). All transmissions are wireless, and node D additionally has a wired Internet connection.

ments, we show performance results when *Turbo-* and *Burst Mode* are enabled and discuss the pros and cons of these modes in access and backhaul networks. The combination of the backbone design with the deployment in a dense urban area with lots of interfering networks makes these measurements unique and important for the understanding of the potential and limitations of wireless backhaul networks. Finally, Section 4 provides background and a discussion comparing *MagNets* to existing related projects, with specific regard to the design, deployment and performance, whereas we draw our conclusions in Section 5.

2. Architecture of the *MagNets* wireless network

The objective of *MagNets* is to deploy a next-generation network testbed that provides novel insight into the use and the behavior of wireless technology [13,19]. The *MagNets* network consists of a high-speed WiFi backbone and an access network equipped of currently planned 100 nodes. All nodes feature high-speed mesh routers, multiple WiFi cards that run in different frequency bands, and some are additionally equipped with heterogeneous technologies, including WiFi, Bluetooth and UMTS. The network is designed as a joint operational-research network, i.e., the network will be accessible as a production network to students at the university, but at the same time we maintain control over the network to perform measurements and deploy and experimentally evaluate new protocols. For the motivation pointed out before, this paper focuses on the backhaul network that has been operational since March 2006.

2.1. Design criteria for backhaul networks

The first step in realizing a high-speed multi-hop backhaul network is a careful engineering to design and plan the network and the nodes. The requirements are thereby fundamentally different from the criteria of the access part where coverage and reachability are the primary objectives.

- **Buildings:** Buildings must be found that (i) provide line of sight for the wireless links, (ii) allow for antenna deployment – technology-wise (e.g., power, Internet connectivity) and administration-wise, (iii) are within wireless transmission distance, (iv) have reasonable one-time installation and recurrent maintenance costs.
- **Topology:** The topology does not need to take a special form, i.e., it can be a linear topology, a tree, a mesh, as long as the links form a coherent network. The topology can be assumed to be static, yet resilience should be taken into account in case of failures.
- **Nodes:** Nodes must primarily be designed for speed, where speed includes the speed of the link to the next neighbor, but also the ability to forward data over multiple hops without performance degradation [17]. The

components used in a backbone node must not be restricted by power or hardware limitations (e.g., antennas).

- **Network management:** A vital part that is often forgotten at this stage is the ability to monitor and manage the network, both for scientific and operational purposes.

2.2. Design of the *MagNets* backbone

Following the above guidelines, we set out to deploy the *MagNets* backbone in practice and in the heart of Berlin, as shown in Fig. 2. It is deployed on five buildings. All premises allowed us to install antennas on their roof, provided power and indoor space to place a router. Moreover, all buildings except *ETF* level well over neighboring buildings and have unobstructed line of sight. The link distances vary between 330 and 920 m, with a total end-to-end distance between T-Labs and T-Systems of 2.3 km. As a particular feature, links 3 and 4 are deployed in parallel to answer the question whether it is possible to double the capacity between two buildings by having two parallel links.

Fig. 3 shows a logical view of the backbone. Two things are important here. First, an out-of-band management network connects the backbone nodes with a central management node. This network, running as tunnels over wires, ensures access to the strategic points of the *MagNets* network. Second, some nodes of the backbone are connected to the mesh cloud that provides access to the users. In the mesh cloud, data may again be forwarded over multiple hops from and to the backbone.

Table 1 summarizes the link characteristics. Of particular interest here is the heterogeneity of the link parameters. Besides the diverse link length, note that the backbone uses both 802.11a and 802.11g. For research purposes, we want to assess how the frequency bands are used in dense urban areas and what impact environmental factors (interference, radar, etc.) have on the backbone.

Each *MagNets* node consists of a Linux PC with a 3 GHz processor and 1 GB of RAM that acts as a router,

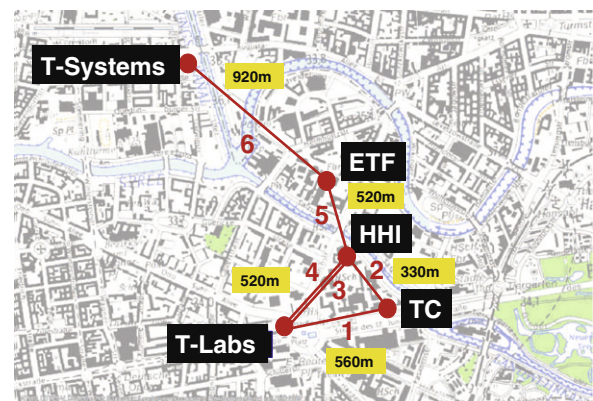


Fig. 2. *MagNets* WiFi backbone in the heart of Berlin. Map: copyright by Google Earth.

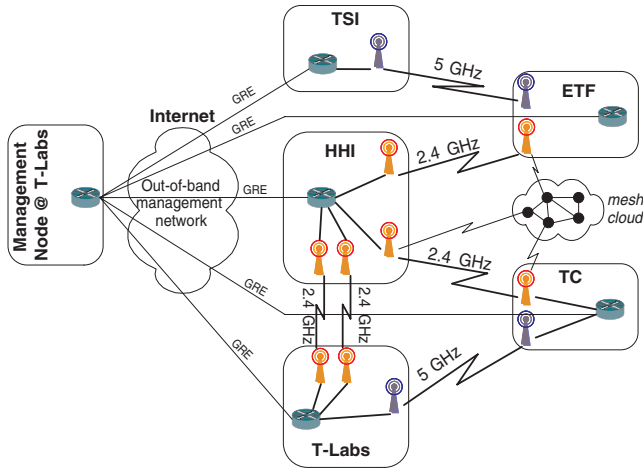


Fig. 3. MagNets layout.

Table 1
Link characteristics

Link	Location	Length (m)	Freq (GHz)
1	TLabs-TC	560	5
2	TC-HHI	330	2.4
3	TLabs-HHI1	520	2.4
4	TLabs-HHI2	520	2.4
5	HHI-ETF	520	2.4
6	ETF-TSI	920	5

as depicted in Fig. 4. We chose a PC-based solution in order to be able to install computationally intensive applications (e.g., network management) as well as application specific services on these nodes (e.g., traffic generators, Web caches).

Attached to the router are one or multiple WiFi access points, one for each outgoing link, on independent network interface cards (HHI, e.g., has a four-port network interface card and four APs attached). The access points are based on Intel IXP420 at 266 MHz (IAP) and IXP425 at 533 MHz (OAP) programmable network processors [6], and they are running a proprietary operating system called Lancom LC.OS [3]. In [7] it has been shown that APs based on the IXP4xx series architecture outperform APs based on Broadcom or IBM processors from the same class. Moreover, all APs feature two proprietary and optional protocols termed *Turbo Mode* and *Burst Mode* that allow rates up to 108 Mbps [4]. An overview of all node components can be found in Table 2.

To limit the damping of the signal on the RF cable, 10 APs are suited for outdoor usage and mounted along the antenna. The backbone uses directional antennas only,

eight operate at 2.4 GHz and the rest at 5 GHz, for two reasons: they are required to bridge the distance between two neighboring APs, and they allow for spatial reuse, i.e., several links can be activated at the same time without causing mutual interference.

In summary, our node setup allows each link to operate independently, i.e., each AP can operate at any time, and thus each node can transmit to each neighbor at the same time. We anticipate here that we therefore overcome the limitations known from early multi-hop networks that only use a single WiFi card [17,18].

2.3. MagNets parameters space

The deployment of the *Magnets* backbone allows us to address the main question raised in this paper, i.e., how far are mesh networks from fulfilling the vision of a ubiquitous high-speed Internet access. This assessment is far from easy, as a plethora of parameters and factors, from temperature over interference of neighboring networks to the interactions of several control loops of the different protocol layers (MAC, TCP, application), influence the end-to-end performance. Therefore, it is important to first lay out the parameter space and define the observation strategy.

The parameter space can be partitioned in three categories: link, topology and traffic parameters, as shown in Table 3. *Link parameters* are those that influence a single link: distance, capacity, frequency, PHY- and MAC-layer protocols such as (*Turbo Mode* and *Burst Mode*). The exploitation of this group of parameters allows, e.g., the systematic evaluation of the MAC-layer bandwidth of the each single link. High end-to-end throughputs can only be expected if the per-link throughput is high.

The second group contains *topology parameters*: where is traffic injected, where is the destination and how many hops does it pass along the way? In the *Magnets* backbone, the number of hops can be easily varied from 1 to 4 and even extended to six if a loop is introduced between T-Labs, HHI and TC. Moreover, topology also includes the interference among the links when multiple links are activated at the same time. We distinguish three different types of interference: (i) interference among links in different frequency bands (2.4 vs. 5 GHz), (ii) interference among links in the same band and (iii) interference between the “twin” links (3 and 4 in Fig. 2).

Finally, the *traffic parameters* captures how the traffic is injected into the network. For measurement purposes, it is vital to assess the throughput under different loads, i.e., as




Fig. 4. MagNets node.

Table 2
Hardware used in the *MagNets* backbone

Component	Vendor	Type	Number	Characteristics
Router	PC	Linux	6	3 GHz, multiple NICs
AP	Lancom	OAP-54	10	54/108 Mb/s
	Lancom	IAP-54	2	54/108 Mb/s
Directional antenna	Lancom	AirLancer Extender O-9a	8	2.4 GHz, 9°/23 dBi
	Wimo	PA13R-18	4	5 GHz, 18 dBi

Table 3
Backbone Parameter space

Group	Parameter	Values
Link	Distance	330–920 m
	Frequency	2.4 and 5 GHz
	Channel	3 and 19 orthogonal channels
	<i>Turbo Mode</i>	On/off
	<i>Burst mode</i>	On/off
Topology	src-dst	Any of the five nodes
	Interference	Single link/all links
	Hop length	1–6 hops
Traffic	Pattern	CBR, VBR
	Packet rate	100–126,500 pps
	Packet size	64–1472 bytes
	Protocol	TCP, UDP

a function of the traffic pattern, the packet rate, the packet size and the higher layer protocols.

The complexity of assessing the parameters arises for two reasons. First, the measurement of a single parameter is already challenging because of the wireless environment. Link rates change as a function of environmental conditions (interference). To perform measurements, it is either necessary to take interference into account (which is a problem in itself again), or to find suited statistical metrics to include but abstract the environmental factors. Second, the parameters have mutual impacts on each other. By changing the channels, the frequency or the modes, the backbone has its impact on environmental factor as well, which then impacts the backbone again. Therefore, a systematic measurement approach is necessary. Our strategy is to vary only one parameter at a time, perform experiments with a duration that is larger than the expected variation time of that parameter and repeat the measurements multiple times. Therefore, a large number of measurements have to be performed to gather a relevant statistical sample space for each aspect to investigate. The presented measurements were taken over multiple months, several measurements a day, resulting in more than 80 GB of data [5].

3. High-speed wireless backbone at work

This section presents experimental results from the *MagNets* backbone. We adhere in each measurement to our methodology to just vary a single parameter at a time. In particular, we search for answers to the following questions:

- What are the characteristics of each link, as a function of the link parameters?
- Is multi-hop transmission the sum of the link transmissions – or (how much) less?
- Can we double the transmission capacity when we use two parallel links?
- Which are the dominant environmental factors that influence the throughput: interference, day and night cycles, or social events such as the *Soccer World Cup* held in Berlin during the summer 2006?

Finally, by aggregating those individual results, we get a good picture of the behavior of the backhaul network.

For the following measurements, we injected traffic from the different routers using *iperf* [2] and *D-ITG* [1]. For UDP traffic, we set the source rate to exceed the link capacity. Measurements typically lasted for 90 s and were repeated 10 times and averaged over intervals of 50 ms.

3.1. Link characteristics

First, we assess the throughput of the individual backbone links. The APs have a raw capacity of 54 Mbps, and deducting roughly 50% overhead from lower layer protocols, we might expect around 27 Mbps at the application layer. But how much throughput do we get in a city like Berlin? Fig. 5 shows our measurement results.

Fig. 5(a) shows the throughput using UDP. The *x*-axis denotes the measured throughput, the *y*-axis denotes the probability density, i.e., a high peak implies that a larger number of samples was measured for this bandwidth. The figure first shows that the link 1 has the highest average throughput (27.8 Mbps) and the lowest variation (2.14 Mbps). The average throughput of links 2–4 varies between 15.9 and 18.7 Mbps. Since the distance of all links is similar, we attribute the superior performance of link 1 to the absence of interference in the 5 GHz band. Finally, links 5 and 6 (not shown in the figure) have a dismal performance with 4.3 and 5.4 Mbps, with very large fluctuations. The reasons for the low throughput are the non-perfect line of sight of the ETF building as well as the higher influence of interference of the lower building.

Fig. 5(b) shows the results for the same experiment except that we use TCP as the underlying transport-layer protocol. The throughput distribution is similar to the UDP measurements, with a slightly lower average. These

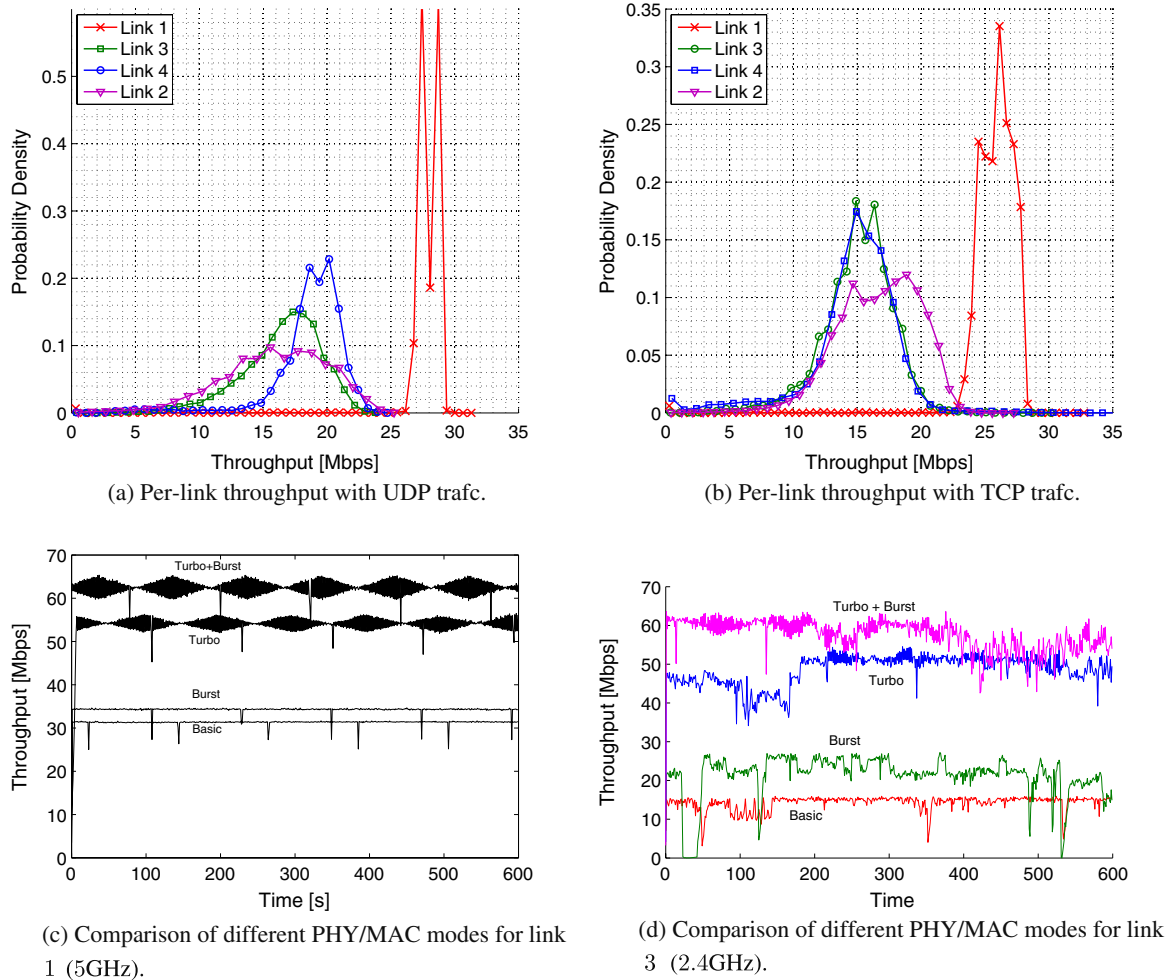


Fig. 5. Link characteristics: (a) PDF of the individual links using UDP, (b) PDF when using TCP, (c) comparison of 802.11, *Turbo*- and *Burst Mode* for link 1 and (d) comparison of 802.11, *Turbo*- and *Burst Mode* for link 3.

results could be expected because TCP reacts to packet loss and the injected traffic therefore does not always saturate the link.

Next, Fig. 5(c) and (d) compare the link characteristics of links 1 and 3, respectively, with basic 802.11a/g, *Turbo Mode* and *Burst Mode* enabled. *Turbo Mode* doubles the transmission rate to 108 Mb/s by enlarging the channel from 20 to 40 MHz. In the 2.4 GHz, the *Turbo Mode* frequency is centered on channel 6 and interferes with all channels in this band. The *Burst Mode* enables an AP to increase its sending rate by waiting only for a shorter SIFS (short inter-frame space) period instead of the standard DIFS (distributed inter-frame space) after a successful transmission.

On link 1, the enhanced modes increase the throughput from 31.3 Mbps with standard 802.11a to 34.2 with *Burst Mode*, 53.8 Mbps with *Turbo Mode* up to 62.4 Mbps when both modes are enabled simultaneously. For link 3, which is in the 2.4 GHz range, the improvements go from 8.4 Mbps to 14.2 (*Burst*), 39.1 (*Turbo*) and 50.3 (both modes). Moreover, the time plots (Fig. 5(c) and (d)) show that the throughput variations increase in absolute terms

with the higher data rates, but remain similar in terms of relative variations.

3.2. Multi-hop

The ability to efficiently relay data from and to the wired Internet is the key task of the backhaul network. However, as pointed out in Section 2, currently deployed backhaul networks have three problems. First, experimental performance analysis shows a dismal performance over multiple hops. While these results point out the limitations of backhaul networks, we ignore their potentials. Second, we ignore the challenges of traffic engineering in wireless backhaul networks. Third, it has been shown that TCP flows severely suffer in multi-hop wireless networks. As we pointed out, the *MagNets* backbone is designed to alleviate this drawback – but to what degree? In the following experiments, we use the topology shown in Fig. 2 but without links 3 and 4, to obtain a linear 4-hop topology from T-Labs to T-Systems. All links are set to standard 802.11a/g.

First, we inject UDP traffic at T-Labs. We create 4 flows, one towards each attached destination, thus creating 1-, 2-,

3- and 4-hop flows. Each flow creates data at 20 Mbps, thus the aggregation of all traffic exceeds the per-link bandwidth measured above. Fig. 6(a) shows the throughput distribution of flows with destination TC and ETF at TC (i.e., after 1-hop) and the throughput of the flow to ETF measured at ETF (3-hop). The other flows and the throughput measured at other locations are not shown. After 1-hop, all flows receive a long-time fair share of 8 Mbps (32 Mbps divided by four flows), even though the short-time throughput may vary considerably among the flows. However, only a fraction of the traffic towards ETF that crosses link 1 eventually reaches ETF due to the low bandwidth on the last hop (2.5 Mbps). That is, a large fraction of packets transmitted on link 1 is dropped at the bottleneck router before link 5.

We can draw two conclusions here. First, the multi-hop throughput corresponds to the per-link throughput of the bottleneck along its path. Thus, we do not see the performance degradation of other, single-AP backhaul networks because *MagNets* uses (i) independent APs and (ii) directional antennas. Second, however, traffic engineering is required to improve the utilization of the backbone. In particular, a significant fraction of packets that cross link 1 are dropped later due to downstream bottlenecks. This issue cannot be resolved with independent transmissions and requires modifications at the ingress to throttle the flows to their downstream bottleneck.

Next, we repeat the same experiment but replaced UDP with TCP traffic. TCP traffic is even more sensitive to multi-hop performance degradation than UDP. However, as Fig. 6(b) shows, the throughput of multi-hop flows corresponds to the throughput of the bottleneck link. Therefore, we can confirm the above conclusion that the *MagNets* backbone defies multi-hop performance drawbacks for both TCP and UDP traffic.

3.3. Twin links

Next, we assess the potential to double the throughput by having two parallel, independent links. In *MagNets*,

links 3 and 4 are deployed as independent links with separate APs and antennas, mounted side by side, running both in the 2.4 GHz band. The distance between the antenna centers is roughly 50 cm and the antenna has a 9° aperture. Is it possible to double the throughput in this way? And if so, what are the link parameters that allow such a doubling? And how does the performance differ from using *Turbo Mode*? What impact does *Burst Mode* have?

Fig. 7 shows the results of this experiment. Fig. 7(a) and (b) show the throughput of links 3 and 4 in the four modes. In basic and *Burst Mode*, the links are set to orthogonal channels (6 and 13); with *Turbo Mode*, the links automatically reset to channel 6. Fig. 7(c) and (d) show time plots of the link activations. We report only UDP measurements, TCP results are similar.

Can we double the throughput with two independent links? Fig. 7(a) and (b) show that we can. First, we measured the link throughput when each link is activated independently. The throughputs reach 8.4 and 6.2 Mbps, respectively. Activating them at the same time doubles the throughput to an average of 14.6 Mbps, thus roughly the sum of the individual link throughputs. Fig. 7(c) shows the throughput over time. Quite interestingly, the links are not activated simultaneously. Instead, one link dominates the transmission for several seconds and then yields to the other link. Thus, the throughput enhancement does not occur because both links are effectively transmitting at the same time, but the links actually interfere and only the “stronger” link transmits, but then at a higher rate.

How does the performance of the twin links compare to *Turbo* and *Burst Mode*? Remember that with *Turbo Mode*, the channel uses twice the channel width, but is set to channel 6 for 802.11g, and the *Burst Mode* reduces the waiting interval after the ack. Fig. 7(d) shows the time plot of the transmissions. As in basic mode, either of the links is transmitting. However, the data rates are higher (due to *Turbo Mode*) and the ramp up is faster (due to *Burst Mode*). As a result, the throughput reaches 23.1 and 22.8 Mbps for the links on average, or 45.8 Mbps in total, when both

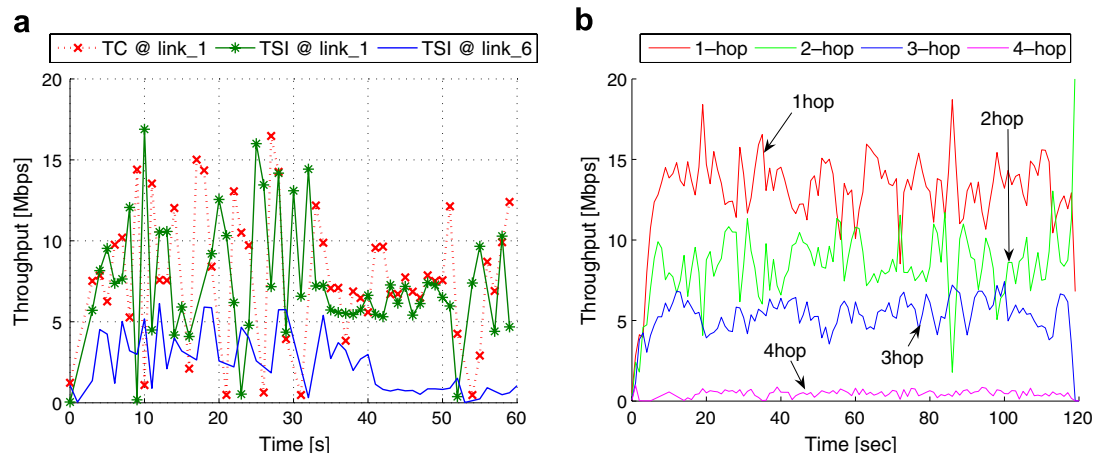


Fig. 6. Multi-hop results. (a) Multi-hop UDP measurements. (b) Multi-hop TCP measurements.

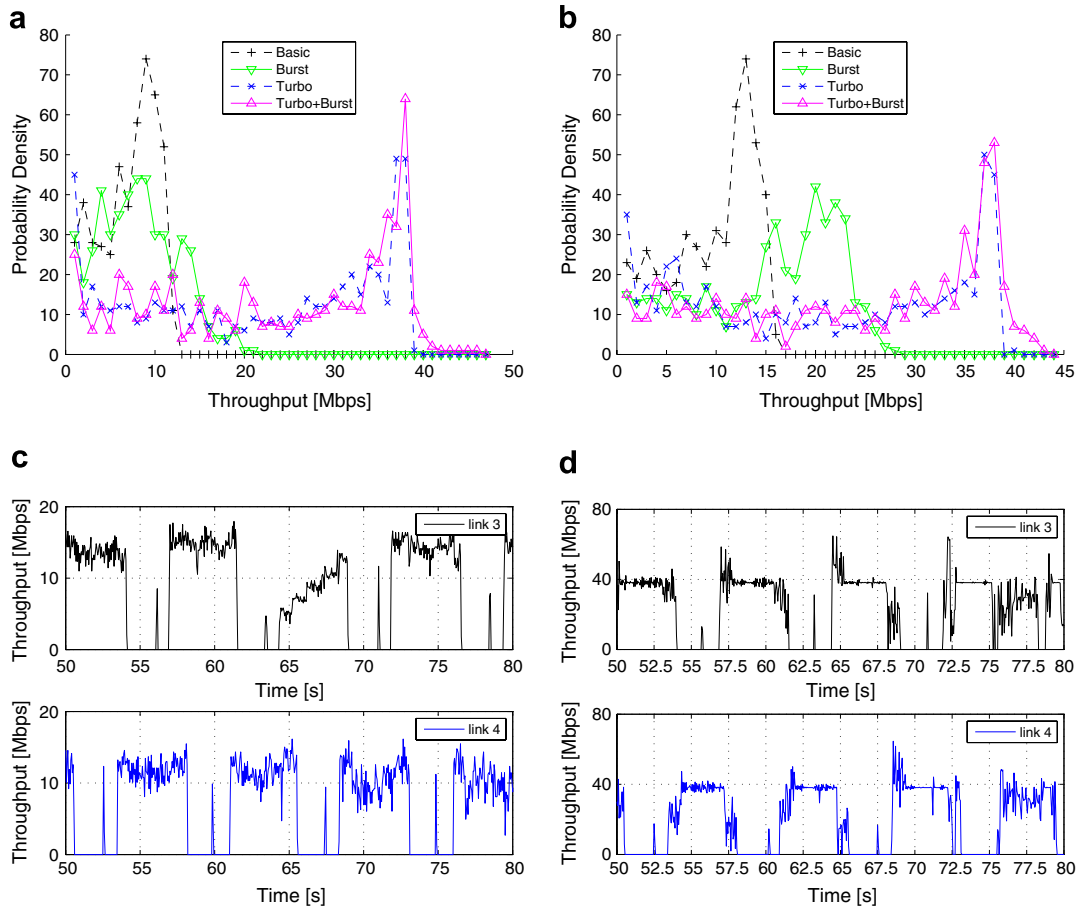


Fig. 7. Can we double the throughput with two parallel links? In basic mode, the answer is yes. With *Turbo Mode*, the answer is no in the 2.4 GHz band due to self-interference. (a) Throughput of link 3 under concurrent activation. (b) Throughput of link 4 under concurrent activation. (c) Twin links in basic mode. (d) Twin links with *Turbo + Burst Mode*.

modes are enabled – a gain by a factor of 3 with respect to the links concurrently operating in normal mode. Therefore, in our measurements, *Turbo* and *Burst Mode* are more efficient than deploying two parallel links. The channel of a single link is sufficiently stable to support a high data rate. The antennas are not sufficiently separated, therefore the link transmissions cause high mutual interference, mainly at receiving side.

3.4. Impact of external factors

What are the dominant factors that influence the throughput of the backbone? We have already reported that the ETF building does not have perfect line of sight because it is lower. As a result, the link at ETF has a significantly reduced throughput. But what other factors that we cannot plan are important? We have assessed three of them here: the impact of interference, the influence of day and night cycles, and the impact of social events such as the World Cup Championship in Berlin in summer 2006.

To measure the impact of interference directly, we lack the ability to record SNRs on the APs because their OS is proprietary. Therefore, we have to deal with a simple

comparison: first, we observe the number of active networks at each node¹ and then we measure the throughput. This methodology is not accurate, but we anticipate that it is sufficient to answer the above questions.

First, we have measured the number of networks that we can scan at each location and on the different channels. The results are interesting: at HHI, in direction T-Labs, we scanned 3, 5 and 2 interfering networks on channel 1, 6 and 11, respectively. At HHI in direction ETF, we found 4, 26 and 11, and at ETF towards HHI we found 2, 14 and 16. The sheer number shows how congested the frequency bands in the 2.4 GHz range are. Moreover, comparing these numbers with the throughput, we see a clear correlation between throughput and the number of competing networks. These results confirm other studies with mesh networks, however, these networks used primarily omni-directional antennas. In contrast, the 5 GHz range is basically still free. This explains why link 1 achieves such high and stable transmission rates.

¹ The APs allow to put the wireless cards into scan mode in order to detect the beacons of other networks. Using this feature, we did multiple scans on all the frequencies before and after the measurements.

Next, we performed measurements over 24 h to assess whether we see differences during the day and the night. Within these 24 h, we measured every 45 min for 120 s. In each measurement, UDP traffic was injected at a rate of 10,000 pkt/s with a payload of 512 bytes, i.e., at rate of 41 Mbps. The resulting traces are sampled at 50 ms. Table 4 shows some statistics obtained on link 3. Besides the mean, medium, maximum, minimum and standard deviation, we also report the *Entropy* and the interquartile range (IQR). The IQR is the difference between the 75th and 25th percentiles. In case of skewed distributions, the median and the IQR provide better insight than the average and the standard deviation, respectively. The *Entropy* of the samples, defined as the *information content* and measured in bits, quantifies the randomness of the considered parameters. To bin the sample distribution we used the Scott Rule [25]. Table 4 shows that during day time, i.e., from 7:45 a.m. to 9:15 p.m., the mean of the throughput was about 1 Mbps lower than in the night (36.85 Mbps compared to 37.90). Moreover, samples collected in the *day* hours are much more spread around their median value (36.78 Mbps). This is witnessed by their higher standard deviation value (3.61 Mbps compared to 2.73 Mbps). Further, the entropy of the *day* samples (4.95 bit) is higher than that of the *night* samples (4.55 bit). All these differences are very likely due to the lower degree of interference during the night. However, the differences are not significant. Thus, we can conclude that the *MagNets* links are slightly influenced by day and night effects.

Finally, we study the impact of special social events, as experienced in Berlin during the 2006 FIFA World Cup. During the games, a large part of the 3.3 million inhabitants of Berlin were watching the game on TV. Therefore some particular activities that can affect the wireless media (e.g., gaming consoles with WiFi connection, mobile phones with WiFi connection, switched-on home WiFi access points) are mostly absent. Do these non-technical variables change the interference patterns or have other effects that may impact the backbone performance? To assess the impact, we performed a set of 14 h long measurements during 5 days on link 3. Eighteen measurements lasting 2 min each were performed with the same parameters described above.

On July 9, the championship final was played in Berlin's Olympiastadion. On July 8, the game was played in Munich, but since the German team played, similar conditions can be expected. As baselines, we measured the parameters on July 6, 7 and 10. As a representative result, in Fig. 8 we show the average packet loss as a function of




Fig. 8. Impact of 2006 FIFA World Cup on Packet Loss of UDP Traffic (link 3).

the day time for the five different days. It is shown that the links are more stable on July 8, 9 from 21:00 to 23:00, i.e., during the matches. However, the differences are not significant and therefore we can state that also these environmental conditions have a negligible effect on the backbone performance.

4. Discussion and related work

Our knowledge about the resource usage (performance, fairness) *in practice* is still in its infancy. Small scale wireless testbeds, e.g., next-generation wireless networks [27] or wireless mesh network testbeds [21], exist in lab environments [26,22]. While they are useful to investigate particular effects, e.g., protocol-specific functions [23,9,21], they lack two vital aspects: scale and interference from other networks. Scale is important for mesh networking research to assess the system's aspect on mesh networks, and interference has been shown to be the primary cause of performance degradations in wireless systems.

Similarly, networks deployed for research purposes often lack the ability to also function as operational networks at the same time [27]. Such networks are of limited use as they may provide novel insight into some protocol behavior under some specifically generated synthetic traffic (often stress tests that show extreme cases), but they do not reflect the behavior under realistic user traffic. Instead, we argue that it is necessary to have both features: the ability to operate under realistic traffic patterns yet the ability to perform measurements and assess protocol behavior.

Table 4
Statistics of the 24 h trace on link 3

	Mean	Min (Mbps)	Max (Mbps)	Median (Mbps)	SD (Mbps)	IQR (Mbps)	Entropy [bit]
Bitrate	37.23	0.00	45.47	36.86	3.36	4.18	5.10
Bitrate day	36.85	0.00	44.24	36.78	3.61	4.26	4.95
Bitrate night	37.90	0.25	45.47	37.19	2.73	4.10	4.55

To this end, *MagNets* has been designed as a joint research-operational network testbed, i.e., the network gives access to university students and allows for experimental deployment and evaluation of network protocols.

Only a few mesh networks today are deployed at scale and provided novel scientific insights. The most prominent network are probably the *MIT roofnet* [12], the TFA network in Houston [15] and the *Digital Gangetic Plains (DGP)* [10,24] in India. Each of these networks has particular features in terms of architecture or environment. In terms of architecture, the MIT roofnet consists mostly of omni-directional antennas, TFA has a mixture of directional and omni-antennas, but all operate in the same frequency band, and DGP uses directional antennas over several kilometers. In terms of environment, the MIT roofnet, e.g., is located in an urban area, TFA is in a sparsely populated area and DGP is in a rural area. Despite its higher scale, *MagNets* has been placed in a dense urban area (Berlin City Center), representing a unique testbed (e.g., in terms of interference) when compared to the networks that are deployed in a rural area of India (*Digital Gangetic Plains (DGP)* [10,24]) and a sparsely populated residential area in Houston, Texas (*TFA* network [15]). When compared to wireless mesh networks deployed in urban areas, *MagNets* presents other unique features: the *MIT roofnet* [12] only contains three-directional antennas and their performance is not evaluated in detail. Moreover, what is in common to the related networks above cited is that the throughput in the backhaul is dismal – typically single-digit throughputs. Therefore, what these networks don't provide is a backhaul network that efficiently connects the APs. In contrast, *MagNets* backbone achieves very high and stable data rates.

Summarizing, *MagNets* presents several positive aspects and unique features in terms of scale, geographical placing, parameters, topology and traffic compared to the above networks. For example, it provides a wide parameters space for investigation: 2.4 and 5 GHz links that span between 330 and 920 m, with the optional enabling of *Turbo* and *Burst Mode* – in contrast, the *MIT roofnet*, *TFA* and *DGP* operate in the 2.4 GHz only. In addition, none of these networks achieve the high rates of *MagNets* as reported in this paper and in [13,19,14]. Indeed we have designed, deployed and evaluated a network that contains a dedicated high-speed backhaul network. It is thereby clear that building such a network has not led to “yet another network” that renders the others useless. Instead, a plethora of parameters, from physical channel selection over interference to higher layer protocols, influence the performance. Each network is unique – this especially holds for wireless environments.

5. Conclusions

How far are mesh networks from fulfilling the vision of a ubiquitous high-speed Internet access? We have noticed that existing wireless backhaul networks that connect the

access points with the Internet are far from achieving the desired throughput. This paper described the deployment of the *MagNets* WiFi backbone and the results we obtained from its experimental performance evaluation.

Our experiences show that a careful planning and deployment of a WiFi backhaul network allow for an application-layer throughput of up to 67 Mbps with off-the-shelf 802.11 Super-A/G equipment and directional antennas. Moreover, having nodes with independent access points that can send and receive data to different neighbors simultaneously does not reduce the throughput over multiple hops. These results clearly contrast previous results, and based on them, we can expect that future 802.11n-based backhaul networks will achieve several hundreds of Mbps throughput.

However, we also experienced severe interference from neighboring networks that reduced the throughput significantly. Especially the 2.4 GHz range shows effects of “pollution” with up to 26 interfering networks on the same channel, whereas the 5 GHz band is (currently) still relatively unused. In contrast, other environmental factors such as day-and-night effects, temperature or social events had little impact on the link characteristics. Therefore, our study confirms that wireless backhaul as a primary solution for delivering broadband services is feasible, also out of metropolitan areas [24], especially to rural communities (where interferences are low).

Thus, to sum up, our final answer is “yes but”: we have shown that, technology-wise, we are on the way to provide ubiquitous high-speed Internet access. However, the increasing pollution of the spectrum imposes a severe threat. Therefore, technical as well as political (spectrum management) progress is vital to successfully deploy wireless mesh networks in the future.

References

- [1] Distributed internet traffic generator, Available from: <http://www.grid.unina.it/software/ITG>.
- [2] Iperf, Available from: <http://dast.nlanr.net/Projects/Iperf>.
- [3] Lancom systems, Available from: <http://www.lancom-systems.de>.
- [4] Lancom systems, 108 mbits super a/g, Available from: <http://www.lancomsystems.de/fileadmin/produkte/features/techpaper/TP-108MbitsEN.pdf>.
- [5] Magnets traces at the traffic web site, Available from: <http://www.grid.unina.it/Traffic/Traces/Magnets.php>.
- [6] Intel IXP4XX Network Processors, Available from: <http://www.intel.com/design/network/products/npfamily/ixp4xx.htm>.
- [7] The tolly group, network processors: Performance analysis of 802.11 broadband routers and access point, Available from: http://www.intel.com/design/network/products/npfamily/tolly_204131.pdf.
- [8] I.F. Akyildiz, X. Wang, A survey on wireless mesh networks, in: IEEE Communications Magazine, vol. 43, no. 9, 2005.
- [9] T. Benzel, B. Braden, D. Kim, C. Neuman, A. Joseph, K. Sklower, R. Ostrenga, S. Schwab, Experience with deter: a testbed for security research, in: Tridentcom 2006, March, 2006.
- [10] P. Bhagwat, B. Raman, D. Sanghi, Turning 802.11 inside-out, in: Proceedings of HotNets-II, Cambridge, MA, November, 2003.
- [11] G. Bianchi, S. Chakraborty, X. Guo, E. Knightly, Multihop wireless mesh networks, in: IEEE Journal on Selected Areas in Communications, vol. 24, no. 11, November, 2006.

- [12] J. Bicket, D. Aguayo, S. Biswas, R. Morris, Architecture and evaluation of the mit roofnet mesh network, in: Proceedings of ACM Mobicom'05, Cologne, Germany, August, 2005.
- [13] A. Botta, I. Matyasovszki, A. Pescapé, R. Karrer, Experimental evaluation and characterization of the magnets wireless backbone, in: Proceedings of ACM Wintech'06, Los Angeles, CA, September, 2006.
- [14] A. Botta, A. Pescapé, G. Ventre, R.P. Karrer, High-speed wireless backbones: measurements from magnets, in: Fourth IEEE International Conference on Broadband Communications, Networks, and Systems (Broadnets), Raileigh, North Carolina, USA, 2007.
- [15] J. Camp, J. Robinson, C. Steger, E. Knightly, Measurement driven deployment of a two-tier urban mesh access network, in: Proceedings of ACM MobiSys, Uppsala, Sweden, 2006.
- [16] Freifunk, Available from: <http://freifunk.net/>.
- [17] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, M. Gerla, The impact of multihop wireless channel on TCP throughput and loss, in: INFOCOM, 2003.
- [18] V. Gambiroza, B. Sadeghi, E. Knightly, End-to-end performance and fairness in multihop wireless backhaul networks, in: Proceedings of IEEE MobiCom, Philadelphia, PA, 2004.
- [19] R. Karrer, I. Matyasovszki, A. Botta, A. Pescapé, Magnets – experiences from deploying a joint research-operational next-generation wireless access network testbed, in: Proceedings of IEEE Tridentcom'07, Orlando, FL, May, 2007.
- [20] R. Karrer, A. Sabharwal, E. Knightly, Enabling large-scale wireless broadband: the case for TAPs, in: Proceedings of HotNets-II, Cambridge, MA, November, 2003.
- [21] E. Nordstrom, P. Gunningberg, H. Lundgren, A testbed and methodology for experimental evaluation of wireless mobile ad hoc networks, in: Tridentcom, February, 2005.
- [22] M. Ott, I. Seskar, R. Siraccusa, M. Singh, Orbit testbed software architecture: supporting experiments as a service, in: Tridentcom, February, 2005.
- [23] J.B. Polglase, C.C. Angelat, C.P. Ruiz, Infrastructure for testing wlan roaming, in: Tridentcom, March, 2006.
- [24] B. Raman, 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, August, 2005.
- [25] D.W. Scott, On optimal and data-based histograms, *Biometrika* 66 (12) (2005) 605–610.
- [26] F. Steuer, M. Elkotob, S. Albayrak, A. Steinbach, Testbed for mobile network operator scenarios, in: Tridentcom, March, 2006.
- [27] M. Takai et al., Scalable testbed for next-generation wireless networking technologies, in: Tridentcom, February, 2005.