MagNets - experiences from deploying a joint research-operational next-generation wireless access network testbed

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

Istvan Matyasovszki
ECE Department, F1-OP-02, FB
University of Limerick
Limerick (Ireland)
istvan.matyasovszki@ul.ie

Alessio Botta and Antonio Pescapé
Dipartimento di Informatica e Sistemistica
University of Napoli "Federico II"
80125 Napoli (Italy)
{a.botta, pescape}@unina.it

***Abstract*—In spite of the recent deployment of wireless access networks, such as meshes and WiFi backbones in cities, the potential and limitations of such networks is still unclear. Deployed networks have a limited ability to gather data or experimentally deploy new protocols, whereas lab testbeds are often limited in scale and lack real applications traffic. This paper presents *MagNets*, a next-generation wireless access network deployed in the city of Berlin. *MagNets* is a joint research-operational testbed that offers connectivity to students, but still allows for experimental deployment of new protocols. We describe the work breakdown and lessons learnt from the design and deployment process. In addition, initial measurement results highlight the potential to shed light on the suitability of wireless technology for next-generation access networks.**

I. Introduction

Wireless technology has the potential to revolutionize society in a way the processor or the Internet did in the last century. Wireless technology will provide ubiquitous and all-time access to an increasing number of devices and foster unforeseeable communication possibilities among humans and machines. A first step towards ubiquitous and all-time communication are wireless access networks that promise to combine the reliability, robustness and wide coverage of cellular networks with high bandwidth known from wireline networks.

Unfortunately, our knowledge of communications over heterogeneous wireless networks is still in its infancy. One of the main reasons is that we lack *semi-productive testbeds*, i.e. testbeds where traffic is created by real users with real applications on the one hand, but where research can be performed, such as deploying new protocols (e.g. MAC or mesh routing protocols) and experimentally evaluate their impact on the user traffic at the same time. Instead, the "testbeds" we are seeing today typically are either operational networks or lab networks. Operational networks, such as the many wireless city mesh networks, carry real traffic, but the access to traffic statistics for analysis is limited and the potential to deploy new protocols is practically zero. In contrast, lab networks used for research purposes provide the flexibility to modify protocols, but are too often toy networks with limited size and no real user traffic.

This paper describes the design and deployment of the *MagNets* testbed [1]. The *MagNets* testbed aims at deploying a next-generation high-speed wireless access infrastructure in the city of Berlin. The network is designed as a wireless access network supported by an operator to perform research, but access is given for free to the students of the Technical University of Berlin to create a semi-productive environment. Moreover, a key feature characteristic of *MagNets* is *heterogeneity* along several dimensions: nodes in the network featuring multiple wireless interfaces with different technologies, such as 802.11, FlashOFDM, 802.16, UMTS and BlueTooth; diverse link characteristics; nodes with varying degrees of processing and storage capabilities; interconnection of multiple mesh networks with disparate routing protocols.

The contributions of this paper are three-fold. First, we describe the design of the *MagNets* network. In particular, we show how the design of the *MagNets* network and its components capture the above objectives. We provide details on the network layout, the details of the network nodes, antennas and masts, the hardware and software choices and the planning time from the initial idea to the final design. Since the planning and deployment of such a network is a complex task, we report on our rationales and processes to build the testbed. We believe that the developed methodologies and processes can be reused for the design of similar types of networks.

Second, we report on the deployment of the testbed. While the deployment basically followed the execution plan, a number of practical issues had to be addressed. We describe the reasons for these problems and how they were solved. The initial plan, together with the lessons learnt, provide fruitful insight that can be used as guidelines for the planning and deployment of other testbeds with similar objectives.

[0] Istvan Matyasovszki was an intern at the Deutsche Telekom Laboratories between March and May 2006.

This work has been partially supported by PRIN 2007 RECIPE Project, by Content EU NoE, and finally by Netqos and OneLab EU projects.

[1] http://www.deutsche-telekom-laboratories.de/∼karrer/magnets.html

Third, we provide initial measurement results. These results highlight the potential to generate novel knowledge about wireless access networks, giving just some examples of the issues that can be investigated using *MagNets*. However, the key challenge is to master the complexity of the results. In particular, the ultimate goal of the network is to provide a reliable and high-quality operation for different types of applications. However, the application-level performance can vary as a function of a plethora of parameters, including physical layer link quality variations, protocol decisions at the MAC, routing or transport layer, and the generated traffic. *MagNets* provides the means to insert the necessary hooks into the infrastructure. However, new tools and methodologies may have to be developed to measure, analyze and understand the network behavior.

This paper is organized as follows. In Section II we describe the design objectives of Magnets and its key distinguishing features compared to other testbeds. Section III describes the planning of the testbed and Section IV describes the deployment of the *MagNets* backbone. Section V presents the measurement strategies and some results to show how *MagNets* can be both studied/analyzed and used for research. Finally, we conclude the paper in Section VI.

II. Testbed objectives

The objective of *MagNets* is to deploy a next-generation wireless access network testbed. The testbed must provide novel insights into the use and the behavior of wireless access networks. In particular, we identify the following requirements for such a testbed

- semi-productive network
- heterogeneous technologies
- strategic and organic deployment
- flexibility and extensibility

First, the network must provide for concurrent research and productive usage. It is imperative that research issues can be pursued, such as the deployment of novel protocols at the MAC and higher layers, after they have been designed and evaluated in simulations. Compared to simulations, the deployment in a productive testbed extends the evaluation of the protocol in a complex environment. However, experimental evaluations are only valid if the network traffic is representative for realistic workloads. Therefore, it is imperative that the testbed is available to a user population. The user population and hence the deployment of the network testbed must be chosen carefully as service quality experienced by the user may vary during operation due to the experimental research properties of the network.

Second, the network must support heterogeneous wireless technologies. A wide variety of wireless technologies exists today, such as GPRS, UMTS, UWB, WiFi, WiMAX, BlueTooth. They feature a wide variety of characteristics, including coverage and capacity. Therefore, they are used for different purposes and can be deployed in parallel. It is important to study multiple technologies in isolation, but also investigate their concurrent use to assess (positive) enhancements and (negative) interference.

Third, *MagNets* must provide relevant output for two communities: research and operators. As outlined above, experimental evaluation of protocols and traffic studies are a fundamental part of research. Unfortunately, the number of testbeds available for such research is limited today, as lab networks are too often limited in size, users, capacity and lack real applications. On the other hand, access to real data from operators is usually hard to get. By contrast, *MagNets* will be able to provide operator grade traffic data to the research community while also yield useful information to operators about capacity constraints, Capex/Opex and eventually even user satisfaction.

Finally, the network must be flexible and extensible in terms of number of nodes, network diameter or capacity. For example, an operator may lay out a network for either high capacity or large coverage. The testbed should be designed to provide both features, e.g. to study the impact and efficiency of a routing protocol on high capacity and sparse density networks.

A. Placing MagNets in literature

Testbeds, e.g. next generation wireless networks [1] or MANET testbeds [2] are often deployed in a lab environment [3], [4] and therefore have a limited scale or, when deployed at large scale, they are placed in a limited area [5]. In other cases, testbeds are targeted toward the study of a specific/narrow research topic [6], [7], [2] and therefore it is difficult to employ them to study a wide range of different topics. Finally, when testbeds are deployed for research purposes, they are often unsuitable to function as operational networks [1] at the same time. Moreover, 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 other relevant networks that are deployed in a rural areas of India (*Digital Gangetic Plains (DGP)* [8] [9]) and a sparsely-populated residential area in Houston, Texas (*TfA* network [10]). When compared to testbed deployed in urban areas, *MagNets* presents other unique features: the *MIT roofnet* [11] only contains 3 directional antennas and their performance is not evaluated in detail. Moreover, *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 m 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. Finally, none of these networks achieve the high rates reported in an initial performance study of the *MagNets* backbone [12].

III. Design

The design of a wireless access network testbed is utterly challenging. A careful planning that takes time, costs and objectives into account is required. Time is particularly crucial

to produce relevant scientific output because a delayed deployment of the testbed may impact the novelty of the results. On the other hand, the deployment of an outdoor testbed is inherently tedious. This section describes our approach to cope with the conflicting issues, by breaking down the network testbed into four independent parts. Then, we describe the individual parts in detail to the degree that they are completed as of now.

A. Project breakdown

The complexity and the challenges of the testbed design and deployment require a phased project breakdown structure. Given the current knowledge of wireless technology, we decided to break down *MagNets* into four phases:

- *MagNets* high-speed WiFi backbone
- *MagNets* WiFi mesh network
- *MagNets* heterogeneous access network
- *MagNets* mesh-of-meshes

All four phases have clear mission statements and objectives. They can be started and executed independently and in parallel. However, in practice, the four phases have different planning and evaluation phases before the deployment, which eventually leads to a staggered deployment. Table I shows an overview of the key planning factors for the 4 types of networks.

The first phase focuses on the deployment of a high-speed WiFi backbone that connects different buildings in the heart of Berlin with off-the-shelf components. Our motivation to deploy the backbone in the first phase is twofold. First, we want to build the backbone with off-the-shelf hardware, so that deployment is easier. Second, we want to exploit the connectivity constraints if data has to be forwarded over multiple high-speed wireless hops. The high-speed backbone enables studies of wireless channel behavior over several hundreds of meters and end-to-end application behavior over multi-hop wireless links. The ultimate insight is whether, and under which conditions, wireless technology can be used to replace wired lines.

The second phase focuses on the deployment of a WiFi mesh network. In contrast to current mesh networks being deployed in various cities, the *MagNets* mesh aims at investigating the limitations in terms of capacity and delay. Since the constraints of current 802.11 protocols are well documented, the mesh network must allow for modifications at the MAC layer. An evaluation of hardware that allows such a customization and is fast and scalable at the same time is required.

In the third phase, the mesh network will be extended with heterogeneous technology to form a heterogeneous 4G network. This phase aims at the use of heterogeneous technology to connect users to the Internet. While WiFi is the most frequently used technology today, the limitation of free spectrum and the limited scalability of a single hot spot cell may give way to the concurrent use of alternative technologies. First, BlueTooth or ZigBee are an alternative to WiFi for low-range communication. Alternatively, some nodes may be equipped with GPRS or UMTS cards. These cards may be used if the multi-hop path through the WiFi mesh is too long for delay sensitive applications such as VoIP. Since GPRS and UMTS networks are already deployed, the testbed will allow us to study how a joint operation of WiFi and alternative technology improves application quality for a user.

Finally, the fourth phase investigates how community networks can be integrated into a single network. It is not uncommon nowadays that multiple mesh networks provide wireless connectivity to isolated “islands” in a city. In Berlin, e.g., two networks besides *MagNets* are already in operation: the first is a community effort named Freifunk.net [2], while another one has been deployed by Humboldt University and is called Berlin Roofnet [3]. Interesting research questions arise when those community networks should be interconnected, as they are operated independently by different administrative authorities. For example, integration of disparate mesh routing protocols with possibly different routing metrics is still an open issue. Specification of policies and their respective effect on inter-mesh routing, as well as mesh gateway functions have not been investigated. The management of a large-scale mesh infrastructure is far from easy and practical experience will prove invaluable on a small scale (within a city), but also on a global scale, e.g. to build a Global Environment for Network Innovations [13].

The following description gives an overview of the current planning stage of *MagNets*. Relating to the 4 phases previously described, the *MagNets* backbone is completed. The mesh nodes have been set up and are currently being evaluated in our Lab. The evaluation thereby includes both WiFi and heterogeneous technology. The inter-connection of multiple community networks is omitted because the project is in its planning phase. Initial locations and hardware are selected, but they have to be evaluated and confirmed.

B. MagNets Backbone Design

The objective of the *MagNets* backbone is to assess if and under which conditions wireless technology can be used to replace wired lines. Our priorities are first and foremost to measure the characteristics of a high-speed wireless network built with off-the-shelf hardware and running current standard protocols at the different layers. By measuring these characteristics, we can identify where shortcomings and bottlenecks occur. Therefore, we identify the following challenges to be addressed in the network and component planning:

- *buildings:* buildings must be found that (i) provide line of sight, (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

[2] http://www.olsrexperiment.de

[3] http://sarwiki.informatik.hu-berlin.de/BerlinRoofNet

TABLE I
GOAL BREAKDOWN FOR THE DIFFERENT NETWORKS AND COMPONENTS.

		Backbone	Mesh	4G	Inter-operation
	Main goal	high speed	high capacity	heterogeneity	protocol boundaries
Network	Building	line of sight	street level	street level	tower
	Topology	linear/tree	mesh	cellular	mesh/linear
	Technology	802.11	802.11	heterogeneous	WiMAX
	Frequency	802.11a/g	802.11a/b/g	lic./unlic.	lic./unlic.
Compo-nent	Antenna	directional	omni / direct.	omni / direct.	omni/direct.
	AP	off-the-shelf	customizable	customizable	vendor-dep.
	Node	off-the-shelf	customizable	customizable	vendor-dep.

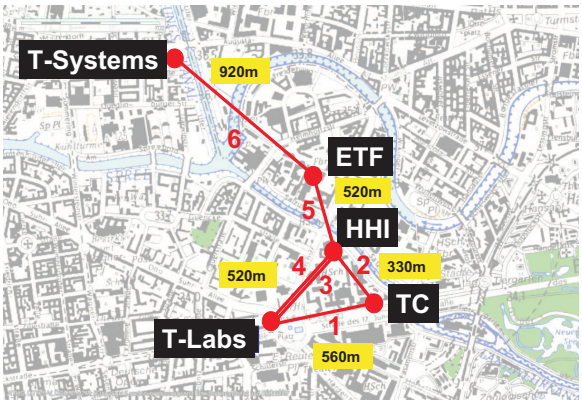

Fig. 1. *MagNets* WiFi backbone in the heart of Berlin.

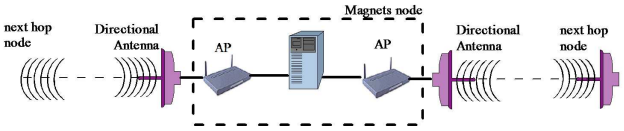

Fig. 2. *MagNets* Node.

long as the links form a coherent network. For research purposes, the topology should be flexible to perform different experiments.

- *nodes:* nodes that are deployed on the buildings must emphasize experimental evaluation, speed and ease of deployment over customizability. Therefore, routers, access points (APs) and antennas may consist of off-the-shelf hardware and may run available software and protocols.

In the subsequent processes of searching suitable locations, we found 5 buildings that suffice our requirements, leading to the backbone topology depicted in Figure 1. Distances between the buildings range from 330m to 920m, with a total end-to-end distance between T-Labs and T-Systems of 2.3km. APs and antennas reside on top of high-rise buildings and have unobstructed line of sight. All transmissions are in the unlicensed spectrum (2.4 and 5 GHz). A *MagNets* node, as depicted in Figure 2, consists of a Linux PC with a 3GHz processor and 1 GB of RAM that acts as a router. Attached to the router are one or multiple WiFi access points, one for each outgoing link (the node at HHI has e.g. 4 APs). The Linux PC is equipped with a corresponding number of network interface cards. Therefore, each link is able to operate independently, i.e. the node is able to perform concurrent transmissions over multiple links.

Table II gives an overview of the hardware used for the *MagNets* backbone. We installed 12 LanCom WiFi APs[4] and 12 directional antennas. To limit the damping of the signal, 10 APs are suited for outdoor usage and mounted along the antenna, to shorten the RF cable length between the antenna and the AP. Only at ETF, indoor APs could be mounted along the inside wall near the antennas. Each AP supports 802.11a/g modes at 54Mb/sec. Moreover, the APs feature two proprietary, optional protocols termed *Turbo Mode* and *Burst Mode*. The *Turbo Mode* doubles the transmission rate to 108Mb/sec 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, in *MagNets*, interference is alleviated because of the directional antennas. In the 5 GHz range, 3 orthogonal Turbo Mode channels are available in the lower band range band and 2 channels in the upper band range. The *Burst Mode* enables an AP to increase its sending rate by waiting only for a shorter SIFS (Short Inter-Frame Space) period after receiving an ack. In contrast, in "normal" mode, the sender has to wait for a Distributed Inter-Frame Space (DIFS) after until it can send a new packet.

The access points are connected to directional antennas, 8 of which operate at 2.4 GHz and the rest at 5 GHz. Thus, while the APs are able to use either 2.4 or 5 GHz, the antennas fix the used spectrum. Directional antennas are required to bridge the distance between two neighboring APs, but also to allow spatial reuse. Since most antennas are mounted on the same pole, directional antennas reduce the interference among the antennas compared to omnidirectional antennas (though some of the main and side lobes may still cause some interference). Finally, directional antennas alleviate unfairness when *Burst Mode* is used: with omnidirectional antennas, the *Burst Mode* may lead to starvation of neighboring senders because it increases the probability that the sender gets access to the channel for the subsequent transmission again.

C. Design of the WiFi Mesh

While the *MagNets* backbone assesses the suitability of wireless technologies in the backhaul, the *MagNets* mesh addresses opportunities and challenges on the last hop. The mesh will shed light on the question how well mesh networks

[4]www.lancom-systems.de

TABLE II
HARDWARE USED IN THE *MagNets* BACKBONE.

component	vendor	type	number	characteristics
Router	PC	Linux	6	3GHz, multiple NICs
AP	Lancom	OAP-54	10	54/108 Mb/sec
AP	Lancom	IAP-54	2	54/108 Mb/sec
dir. antenna	Lancom	AirLancer Extender O-9a	8	2.4GHz, 9°/23 dBi
dir. antenna	Wimo	PA13R-18	4	5GHz, 18dBi

and multi-radio technologies scale in terms of capacity and connectivity. Current 802.11 technology used in HotSpots achieves rates up to 108 Mb/sec. However, for the demand of future networks, as e.g. outlined in the 100x100 project [14], the capacity must scale up to several 100s of Mb/sec or even Gb/sec. We therefore investigate up to which scale such a high capacity can be achieved by careful capacity planning and how well the network can sustain delay- and capacity-sensitive applications. Towards this objective, we identify the following challenges to be addressed in the network and component planning:

- *nodes:* every single node must support transmissions of at least 100 Mb/sec.
- *network planning:* the capacity in the entire network coverage area must support at least 100 Mb/sec.
- *protocols:* if necessary, novel protocols must be developed that ensure an efficient usage of the capacity.

Our approach to address the WiFi mesh requirements is to first select and evaluate mesh nodes that have the potential to achieve 100 Mb/sec transmission speed. There are two possible options: build customized high-speed hardware or combine existing components to boost the transmission rates. We have waived the approach to build customized hardware because vendors have a greater potential to develop such hardware. Instead, we aim at scaling the network capacity by adding multiple WiFi cards into a single node. In particular, we have chosen two pieces of hardware for evaluation: routerboards and mini-ITX boards.

The RouterBoard 532 [5] series are an all-in-one integrated communication platform. It features a MIPS32 CPU running at up to 400 MHz and a 32-bit PCI controller at 66 MHz. For networking, the board provides up to 3 Ethernet ports and 2 MiniPCI slots on board. Daughterboards can additionally be attached via on-board connectors. The RouterBoard 564, e.g., is a daughterboard that provides 6 Ethernet ports and 4 MiniPCI slots. Using Atheros [6] 802.11a/g WiFi cards that offer 54 Mb/sec in their standard mode and 108 Mb/sec with SuperAG enhanced technology, the theoretical throughput of a routerboard reaches up to 648 Mb/sec. Mini-ITX boards [7] are small-scale but fully equipped PCs. Thus, Mini-ITX boards come in a large variety of processor power, RAM and bus speeds. Mini-PCI WiFi cards can be connected via 1- or 2-slot PCI Riser Cards and Mini-PCI to PCI cards.

Comparing the current availability of hardware, we find that the RouterBoards have their main advantage in that cases are available for indoor and outdoor use, whereas it is difficult to find outdoor-proof cases for Mini-ITX boards. In contrast, the Mini-ITX boards excel in their flexibility to choose CPU and memory. Moreover, Routerboards fail to provide USB ports, which limit the direct attachment of additional hardware (e.g. WebCams, storage). The price for both types of boards with the option to hold 6 Mini-PCI WiFi cards is currently between 250 and 300 US$. Moreover, both boards support open source software (Linux, MadWiFi, etc). The availability of customizable tools ensures that *MagNets* provides ample opportunities and flexibility to deploy and evaluate protocols at any layer. It will allow experimental evaluation of cross-layer optimizations that have been proposed in the research literature [15] and to shed practical, experimental light on the benefits and drawbacks of wireless access networks.

[5] www.routerboard.com
[6] www.atheros.com
[7] www.mini-itx.com

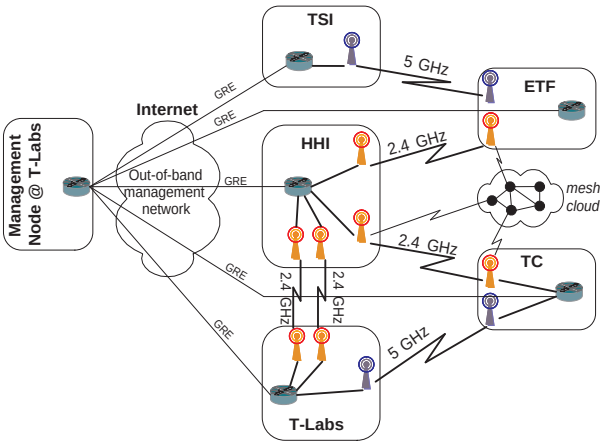

Fig. 3. *MagNets* backbone network structure

D. 4G network

Fourth-generation (4G-) networks focus on the use of heterogeneous technologies within the same network, such as WiFi, UMTS, WiMAX, etc. The current availability of GPRS and UMTS networks in Berlin allows the integration of corresponding cards into the *MagNets* nodes. By super-imposing multiple network configurations on *MagNets*, issues such as TCP performance during vertical handovers between multiple access technologies can be explored. Of particular interest are operator-driven optimizations for resource management and load balancing, as well as opportunities for separation of control and data planes that exploit diverse characteristics of wireless technologies.

IV. DEPLOYMENT

In this section, we report our experiences from deploying the *MagNets* backbone, describe practical issues that had to be addressed and report the lessons learnt.

A. Hardware deployment

All backbone routers have the exact same hardware configuration except the number of NICs. Therefore, we automated

the initial setup of these nodes using the Linux Disk Dump[8] (*dd*) utility. It proved to be a fast and effective approach to clone Fedora Core 4 with a customized 2.6 series kernel to all nodes. Therefore, we plan to use the same approach for the mesh nodes, i.e. to clone the OS and a basic configuration skeleton onto the flash memory cards of the mesh nodes.

B. Management network

Each *MagNets* node contains one additional NIC that connects the router to an out-of-band management network shown in Figure 3. The management network has a number of tasks: it facilitates AP firmware updates, backbone router OS updates, changes to routing configurations and protocols, AP and router configuration backup, log file transfer for centralized processing, debugging of wireless APs and links, time synchronization, traffic trace collection and SNMP statistics monitoring.

The connection from each backbone node to the main management node located at Tlabs is tunneled through the public Internet using authenticated Generic Routing Encapsulation (GRE) tunnels. The GRE tunnels provide virtual connections among the backbone routers and the main network management node. The management network is not included in the overall IP routing and addressing scheme of the *MagNets* backbone to avoid that routing forwards traffic from the backbone via the management network to the Internet. Although GRE does not provide encryption it has been chosen for its simplicity, performance and wide support in the various networks of the organizations/companies hosting the backbone routers. To enhance security, the management links will be migrated to IPSec after availability is tested at all locations.

The management network also provides time synchronization for the backbone routers, which in turn synchronize their associated wireless APs and later mesh nodes. In the current deployment, the Network Time Protocol (NTP) is used for time synchronization. However, the clock skew caused by short delays and high bandwidth make it difficult to correctly measure and interpret network parameters, so that a GPS based solution is being prepared for future deployment.

C. Backbone Routing

Addresses in the *MagNets* backbone use a set of private IPv4 address spaces (subnets). The backbone supports both static and dynamic routing. Static routing allows the backbone topology to be shaped according to specific needs. Moreover, static routing reduces the impact of the network layer on end-to-end throughput. In contrast, dynamic routing allows deploying and evaluating different routing protocols. In the backbone, due to topology characteristics, OSPF is the dynamic routing protocol of choice, as OSPF allows variable size sub-netting, has a low traffic overhead and supports authentication. Using OSPF, the backbone already provides the potential to study traffic load balancing on the parallel links between T-Labs and HHI.

[8]http://www.gnu.org

D. Measurement/monitoring setup

Network measurement, traffic generation and traffic trace collection are simplified by the management network. The central management node at T-Labs is running an SNMP based network statistics monitoring tool called Cacti[9] that periodically (every 5 minutes) polls backbone routers and APs to generate long term disk and network interface utilization graphs.

To generate synthetic traffic on the backbone, we use two tools: Iperf[10] and D-ITG[11]. D-ITG (Distributed Internet Traffic Generator) is a tool capable of generating traffic accurately replicating appropriate stochastic processes for both IDT (Inter Departure Time) and PS (Packet Size) random variables (e.g. exponential, uniform, cauchy, normal, pareto). D-ITG supports both IPv4 and IPv6 traffic generation and it is capable of generating traffic at network, transport, and application layer.

Backbone routers have the capability to capture, process and archive traffic traces at layer 3 and above, but cannot capture layer 2 traces as they are not directly connected to wireless NICs inside the AP. However, collecting layer 2 data is important for the characterization of the wireless links. This lack is a clear drawback of the off-the-shelf setup, even though a limited amount of layer 2 information can be collected from the APs. However, the output can only be sent to the standard output. Therefore, additional effort is needed to gather the traces on the Linux router and to correlate this information to higher-layer information. The mesh nodes being deployed in the second phase of the project will make use of the monitor mode of their wireless NICs to capture 802.11 frame information including the IEEE 802.11 header as well as physical layer information.

E. Interference

The *MagNets* backbone features 802.11a and 802.11g links. Since the antennas only support either version, the technology has to be fixed for each link. For the deployment, we have chosen the configuration that allowed for most variation to study wireless channel characteristics. Even though the attenuation is higher for 5 GHz, regulations allows higher transmission power for 802.11a. Therefore, we have selected links 1 and 6 to be 802.11a. This configuration allows a comparison of a 920 m link with a 560 m link. Moreover, we can compare links 1 and 3, which both span 500+ m but use 802.11a and 802.11g respectively. Similarly, we can compare 802.11g links over 330 m and 520 m.

Interference is generally higher in the 2.4 GHz range due to competing wireless networks. Cordless phones, microwave ovens, and Bluetooth devices are also common in the 2.4 GHz band. While their interference is negligible for point to point backbone links, they will have to be considered for the mesh networks. For the MagNets backbone, we observed a varying number of competing wireless networks from different

[9]http://cacti.net
[10]http://dast.nlanr.net/Projects/Iperf/
[11]http://www.grid.unina.it/software/ITG

TABLE III
NUMBER OF COMPETING NETWORKS AT ETF AND HHI NODES.

Location	Direction	Channel			
		1	6	11	13
HHI	TLabs	3	5	2	2
HHI	ETF	4	26	11	
ETF	HHI	2	14	16	

institutions, homes and commercial networks. Table III gives an overview of the number of competing networks for the most affected locations at HHI and ETF. At ETF, between 13 and 16 competing networks were detected at channel 11, while at HHI 26 competing networks were observed on channel 6. At the other backbone node location the number of competing networks is around 3. Furthermore, the competing APs at ETF are mounted relatively close to the *MagNets* APs. Therefore, performance loss in that particular case can also be attributed to near field effect caused by all APs (including *MagNets*) in the vicinity. Moreover, the *MagNets* links cause interference among themselves: at HHI and T-Labs, 3 directional antennas are mounted side by side. While the Freznel zones are comparably small (TLabs-HHI links have a Fresnel *zone 1* radius of 4 m, TLabs-TC 2.64 m, TC-HHI 3.2 m, HHI-ETF 4 m and ETF-TSystems 3.4 m), a detailed evaluation is needed to assess the impact of side lobes and inter-link interference on the link performance.

Less interference is expected in the 5 GHz range since only few networks are using this frequency and because there are 24 non-overlapping channels in the 5GHz band. However, according to official EU documents[12] in several member states, the operation of military and meteorological radars takes place in bands between 5.25 GHz and 5.85 GHz. In fact, in the initial deployment of *MagNets*, we measured strong pulses that probably originate from airport radars. These pulses require the APs to find and synchronize on a new channel - causing several seconds of transmission interruptions. To protect against these sources of interference, APs use two mechanisms: dynamic frequency selection (DFS) and transmit power control (TPC). DFS aims at avoiding interference while TPC adjusts the transmission power to the minimum necessary for a given communication to avoid interfering with radars. Initial channel selection is done based on country-settings, while subsequent channel selection is based on a repeated scanning procedure until a channel without radar signal interference and as few as possible competing networks is identified.

F. Main lessons learnt

The deployment of the *MagNets* backbone, from its first idea to the first bit transmitted, took almost one year. The time was almost equally divided into planning and deployment. It turned out that most of the ideas and visions could be realized in the backbone. Given the objectives, we were able to build a network where fast deployment, flexibility, and ease of use is ensured via off-the-shelf hardware and open source software. The backbone configuration with fast cloning procedures and the out-of-band management network have proven to be very effective. The current drawback of our design is the lack of fully functional layer 2 traffic capture.

In spite of the overall success, a number of details did not turn out as planned. First, it is far from easy to find suited buildings. In particular, roofs are increasingly crowded with antennas that block the view or cause interference, and nearby rooms to host a PC with Internet access are difficult to find.

Second, not all of the backbone nodes achieve the high throughput we hoped for. Deviations from a perfect line of sight, interferences from competing networks and self-interference result in lower throughput. Moreover, we ignored the existence and impact of radar interference. Even though we tried to take all factors during the design into account, their quantitative impact can not easily be assessed.

Finally, at this stage, it was still unclear to us what kind of performance we can expect from the backbone in practise. Optimistically thinking, the backbone has been so well engineered that we should expect an end-to-end transport-layer throughput of roughly half the physical capacity - i.e. 27 Mbps via standard 802.11 and maybe 54 Mbps using *Turbo Mode* and *Burst Mode*. On the other hand, we completely ignore the impact of the large unknowns, in particular interference, but also distance. To which degree do they affect the throughput? Will the throughput degradation be similar for all links? Which parameters are the dominant parameters, or is only a combination of parameters responsible for significant throughput degradations? Are the effects constant or do they incur variations - if so, at which time scale? The following section addresses these questions. The results will not be utterly surprising in the sense that we will not measure throughputs beyond the nominal capacity, but they will yield vital insight into today's reality. We argue, however, that exactly such numbers are important because they are not available in todays simulations and therefore can be used in future models and future algorithms for wireless networks.

V. MEASURING THE *MagNets* BACKBONE

The *MagNets* testbed allows a broad range of parameters to be set and therefore requires a wide range of different measurements to be performed, even if we restrict our attention "only" to the backbone. The goal of this section is neither to provide a comprehensive analysis of Magnets performance nor to deeply investigate some particular aspect of it. Instead, we highlight the range of parameters that can be measured in order to make the reader perceive the potential of such a testbed.

A. Parameter Space

The parameter space can be partitioned in three categories: link, topology, and traffic parameters respectively, as shown in Table IV.

Link parameters capture the parameters that influence a single link, such as distance, capacity, and frequency. In particular, capacity enhancements are achieved using the (*Turbo Mode* and *Burst Mode*) or using 2 parallel links on orthogonal channels. The exploitation of this group of parameters allows the systematic evaluation of the MAC-layer bandwidth of

[12] 2005/513/EC

TABLE IV
BACKBONE PARAMETER SPACE.

Level	Parameter	Values
Link	Distance	330 - 920 m
	Frequency	2.4 and 5 GHz
	Channel	3 and 19 orthogonal channels
	Turbo Mode	on/off
	Burst mode	on/off
Topology	src-dst	any of the 5 nodes
	interference	single link / all links
	hop length	1 – 6 hops
Traffic	Pattern	CBR, VBR
	Packet Rate	100 - 126500 pps
	Packet size	64 - 1472 Bytes
	Protocol	TCP, UDP

the each single link as well as of the end-to-end throughput over large time scales. This investigation provides fundamental insights into the suitability of 802.11 backhaul networks, in contrast to the use of 802.11 with omnidirectional antennas for mesh networks.

The second group contains topology parameters. Alternative network topologies can be created by choosing different traffic sources and destinations. Moreover, the path length can be varied from 1 to 4 hops. It can even be increased to 6 introducing a loop between T-Labs, HHI, TC and back to T-Labs and configuring accordingly the routing on both links between T-Labs and HHI. The topology also captures the ability to activate multiple links simultaneously. On one hand, such an activation increases the number of transmissions, but it also may imply interference. Will such an activation be beneficial for end-to-end throughput or will it hamper it? Will it improve or reduce the end-to-end throughput of TCP over multiple wireless hops, which we know is very sensitive [16].

As for the last group, the traffic injected into the network allows an assessment of different sensitivity parameters of the backbone as well as the capability of the network to effectively transport peculiar classes of traffic (e.g. real time). In particular, several constant and variable (random or non-random) patterns can be used to profile the measurement traffic injected into the network, and the sending rate, the packet size, and the protocol (TCP and UDP) can be varied.

B. Measurement approach

To cope with such a wide parameter space, a systematic measurement approach is necessary. For this reason, before starting to measure the backbone, an accurate planning activity has been conducted to derive the measurement strategy. The main goal of our approach is the correct identification of the responsible parameters of the experimented performance. To achieve this goal, we carefully avoid the simultaneous variation of more than one controllable parameter, i.e. in each measurement stage, just one parameter is tuned. Due to environmental factors, a large number of measurements has to be performed to gather a relevant statistical sample space for each aspect to investigate. Even if this aspect can be seen as a methodology drawback, a fine tuning capability allows to perform very detailed measurements and to address issues difficult to investigate with other testbeds.

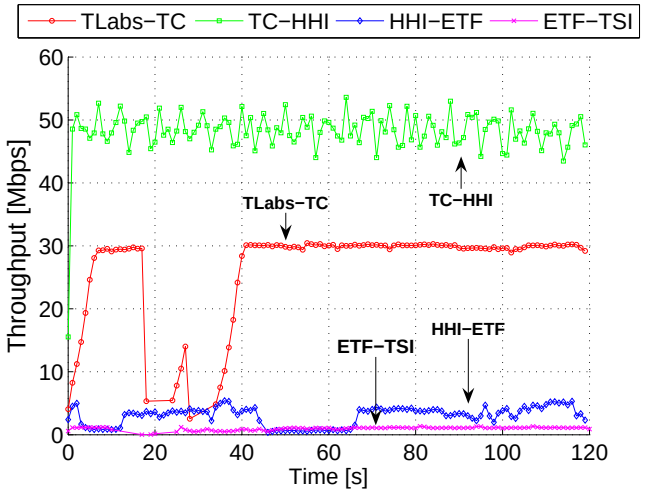

Fig. 4. Per-link characteristics of the Magnets backbone.

C. Results

The following Sections show the capability of *MagNets* to test specific aspects of a wireless network setup. In particular, for each of the three domains of the parameter space (see Table IV), we show some preliminary results. We thereby show that the testbed is able to characterize various aspects of wireless communication and not just suited to study a single parameter.

1) Impact of Link Characteristics: Figure 4 shows the throughput of 4 backbone links. The x-axis denotes the measurement time in seconds, the y-axis shows the UDP throughput in Mbps. The throughput was measured on one link at a time to avoid interference among the links. While the measurements just show the throughput at a specific point in time, the main characteristics of the links have been observed in multiple time intervals. This measurement already hints at many interesting aspects, such as short- and long-term link variation. We note, e.g., that links 2 and 5 show throughput variations in the order of 10% around the average value, whereas links 1 and 6 show little short-term variations. We attribute the absence of variation at links 1 and 6 to the absense of interfering networks in the 5 GHz range and to the dynamic power selection that is only available for 802.11a (5GHz). Focusing on the throughput differences among the links, link 2 (TC-HHI, with Turbo- and Burst-mode enabled) achieves 48 Mbps, link 1 (TLabs-TC) 28 Mbps, link 5 (HHI-ETF) 6 Mbps and link 6 (ETF-TSI) 2.5 Mbps. The main reason for the low throughput of links 5 and 6 is that the ETF building is not as high as the others. Surrounding buildings, obstructions in the line of sight, the length of link 6 and an increase in interference from neighboring APs are responsible for the low throughput. For any network deployed in reality, these or similar factors may have an impact on the operational challenges.

2) Topology - Impact of Multi-hop: Here, we present initial measurements on the *MagNets* backbone that point at issues that have to be addressed for multi-hop wireless networks. In particular, we show measurements that emphasize the need

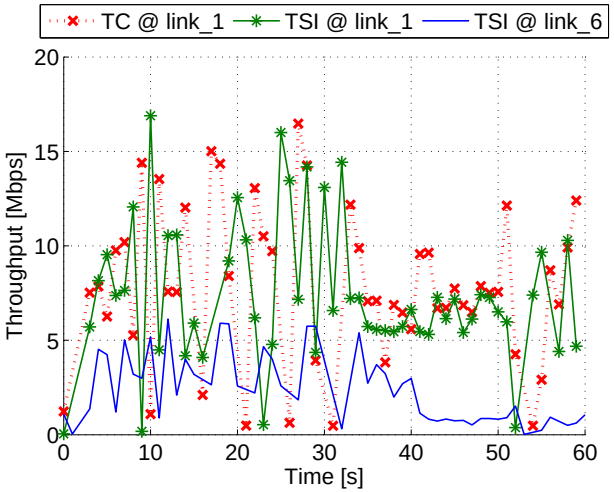

Fig. 5. Multi-hop UDP measurements.

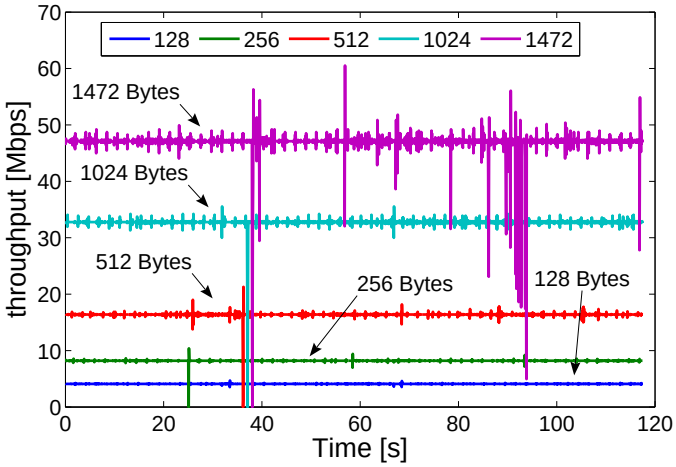

Fig. 6. UDP throughput with different PS and a packet rate of 4000 pps.

for network-wide traffic control to ensure a fair and efficient resource usage in wireless multi-hop networks. While previous work has primarily focused on highlighting fairness issues using simulations, we are the first to emphasize the need for enhanced traffic control in high-speed wireless networks. In the following experiments, we use the topology shown in Figure 1 but without links 3 and 4. The resulting topology is linear, with a maximum of 4 hops.

Consider the objective to optimize the throughput along the backbone. Assume that traffic enters the backbone from the Internet at TLabs and users are attached at any other backbone node. We inject UDP traffic at a rate of 20 Mbps each towards each destination node. Figure 5 shows the throughput of flows with destination TC and TSystems (TSI) at link 1 and the throughput of the flow to TSI measured at TSI (link 6). The other flows (to HHI and ETF) and the throughput measured at other locations is not shown. At link 1, all flows receive a long-time fair share of 12 Mbps (48 Mbps divided by 4 flows), even though the short-time throughput may vary considerably among the flows. However, only a fraction of the traffic (2.5 Mbps) destined for TSI eventually reaches TSI due to the low bandwidth on the last hop. That is, a large fraction of packets transmitted on link 1 is dropped at the bottleneck router before link 6. To achieve an efficient usage of the network-wide capacity, i.e. to avoid that the bandwidth over the first hops is wasted for packets that are eventually dropped anyway, the bandwidth of each flow should be throttled to the bottleneck capacity along its path at the ingress node of the backbone.

3) Traffic - Impact of IDT/PS: In this section we aim at understanding the impact of different combinations of Packet Size (PS) and Inter Departure Time (IDT) on *Magnets* throughput. To clearly understand each single contribution, we first study the impact of PS and after we analyze the joint impact of PS and IDT.

To study the sensitivity of the network throughput to the PS, we injected a given number of packets per seconds into the network but with different packet sizes. Figure 6 shows the UDP throughput (y-axis) as a function of the measurement time (x-axis) achieved sending 4000 pps with packet sizes ranging from 128 to 1472 Bytes. Despite the obvious decrease in throughput, we note that the plots become more irregular as the packet size increases. In particular, small PS (128 and 256 Bytes) result in an almost straight line. High PS (1024 and 1472 Bytes) result in high-frequency oscillations with spikes that drop even down to 0 Mbps. This behavior is due to the increase of the imposed bitrate. In fact, thanks to point-to-point configuration of its links, the *MagNets* backbone transports packets of all the sizes we tested with equal performance. This result is true for specific packet rates. Increasing the packet rates above 6000 pps, yields, for certain PS, a decrease in the overall performance as the packet loss becomes significant.

To study the sensitivity of the network throughput to combinations of PS and IDT, we generated six interleaved flows with the same imposed throughput but achieved with different combinations of PS and IDT. The flows were characterized by PSs ranging from 64 to 1472 Bytes and IDTs ranging from 2750 to 126500 pps. The test was repeated twice: first, PS/IDT pairs are chosen such that the imposed throughput is about 65 Mbps (saturated link); then, a combination of PS and IDT is selected that produces an imposed throughput of about 33 Mbps (far from saturation). Because all the flows demand the same throughput, this configuration allows to understand the impact of the packet rates and sizes on *Magnets*.

Figure 7 depicts the PDF of the throughput samples for the two imposed throughputs. Inside the figure, each flow is labeled with the PS (first number) and IDT (second number). The different pairs of IDT/PS result in significantly different throughputs. We attribute this difference to the packet rate used by the flows with a PS lower than 512 Bytes, and only little to the overhead of low-layer headers. Next, we notice that the flows characterized by packet rates higher than 63250 pps are not able to generate any packet independent of the imposed throughput. Finally, we conclude that throughput improvements are easier by increasing the packet rate rather than using a large PS, independent of the imposed load (33 Mbps and 65 Mbps).

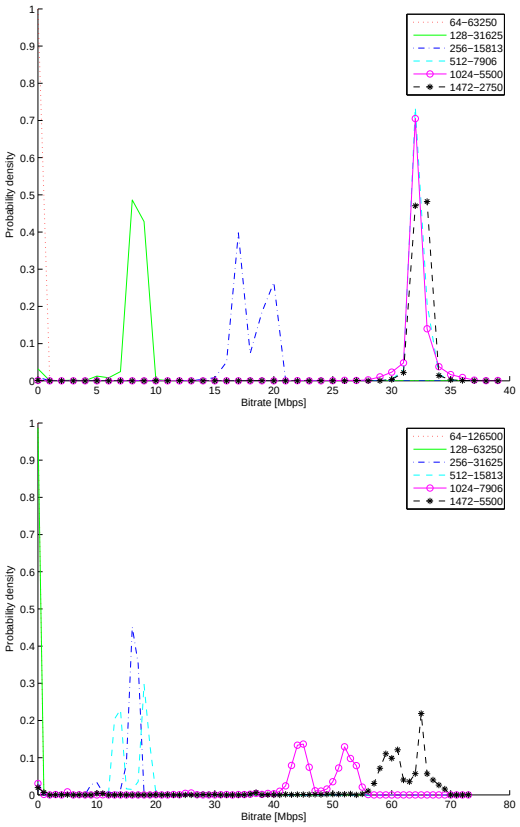

Fig. 7. Throughput of different PS/IDT with the same imposed value.

D. Discussion

These results emphasize that a wireless access network is more than the sum of its nodes. *MagNets* allows us to quantify the impact of wireless technology on efficiency and fairness in a high-speed wireless network. While the initial data was gathered with directional antennas, more data using different configurations, such as omnidirectional antennas, will be gathered when the mesh nodes are deployed. Having pointed out the importance of these issues, researchers can now devise new protocols for wireless access networks that increase efficiency and mitigate unfairness. After the protocol design, implementation and simulations, *MagNets* will fulfill its second objective: to be an evaluation testbed for new protocols. The testbed will finally complement the knowledge gained from simulations.

VI. Conclusions

MagNets is a next-generation wireless access network testbed. The testbed is designed to deploy novel protocols and study their characteristics, such as performance or robustness, in a realistic telecoms production environment. In addition, the testbed contains components that can dynamically be added or removed to form different network topologies and to forward real user traffic. This joint research and operation deployment will shed light on the potential and limitations of future wireless access networks.

The design of an access network that integrates multiple wireless technologies, including WiFi, WiMAX, UMTS and BlueTooth and that provides transmission speeds and capacities of several 100 Mb/sec in a densely populated city area requires careful planning. The deployment of the *MagNets* backbone and the initial measurements confirm the successful work breakdown and planning that can be reused for related testbeds.

The lessons learnt from the testbed deployment are vital for upcoming future network initiatives, such as GENI and FIND. *MagNets*' flexibility to experimentally test and evaluate new protocols at multiple layers for multiple technologies provide the fundamental to contribute future network research.

Acknowledgments

We would like to thank Marco Pfeuffer and Wolfgang Kreher from T-Systems International GmbH Systems Integration for helping with the planning and configuration of the Magnets backbone, and Bernd Kanwischer and his team from Vivento Technical Services GmbH for setting up the masts and antennas.

References

[1] M. Takai and et al., "Scalable testbed for next-generation wireless networking technologies," in *Tridentcom 2005*, Feb. 2005.
[2] E. Nordstrom, P. Gunningberg, and H. Lundgren, "A testbed and methodology for experimental evaluation of wireless mobile ad hoc networks," in *Tridentcom 2005*, Feb. 2005.
[3] F. Steuer, M. Elkotob, S. Albayrak, and A. Steinbach, "Testbed for mobile network operator scenarios," in *Tridentcom 2006*, Mar. 2006.
[4] M. Ott, I. Seskar, R. Siraccusa, and M. Singh, "Orbit testbed software architecture: Supporting experiments as a service," in *Tridentcom 2005*, Feb. 2005.
[5] X. Zhong, H.-H. Chan, T. J. Rogers, C. P. Rosenberg, and E. J. Coyle, "The development and estadium testbeds for research and development of wireless services for large-scale sports venues," in *Tridentcom 2006*, Mar. 2006.
[6] J. B. Polglase, C. C. Angelat, and C. P. Ruiz, "Infrastructure for testing wlan roaming," in *Tridentcom 2006*, Mar. 2006.
[7] T. Benzel, B. Braden, D. Kim, C. Neuman, A. Joseph, K. Sklower, R. Ostrenga, and S. Schwab, "Experience with deter: A testbed for security research," in *Tridentcom 2006*, Mar. 2006.
[8] P. Bhagwat, B. Raman, and D. Sanghi, "Turning 802.11 inside-out," in *Proceedings of HotNets-II*, (Cambridge, MA), Nov. 2003.
[9] 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.
[10] 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.
[11] 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.
[12] A. Botta, I. Matyasovszki, A. Pescapé, and R. Karrer, "Experimental Evaluation and Characterization of the Magnets Wireless Backbone," in *Proceedings of ACM Wintech'06*, (Los Angeles, CA), Sept. 2006.
[13] National Science Foundation, "Geni: Global environment for network innovations - conceptual design and project execution plan," Jan. 2006.
[14] "100x100 clean slate project." http://100x100network.org/.
[15] V. Srivastava and M. Motani, "Cross-layer design: a survey and the road ahead," *IEEE Communications Magazine*, vol. 43, p. 112119, Dec. 2005.
[16] Z. Fu, P. Zerfos, H. Luo, S. Lu, L. Zhang, and M. Gerla, "The impact of multihop wireless channel on TCP throughput and loss.," in *INFOCOM*, 2003.