Hindawi Publishing Corporation
EURASIP Journal on Wireless Communications and Networking
Volume 2008, Article ID 274790, 10 pages
doi:10.1155/2008/274790

Research Article

Challenges in Second-Generation Wireless Mesh Networks

Roger P. Karrer,[1] Antonio Pescapé,[2] and Thomas Huehn[1]

[1] *Deutsche Telekom Laboratories, Technical University of Berlin, 10587 Berlin, Germany*
[2] *Dipartimento di Informatica e Sistemistica, Università degli Studi di Napoli Federico II, 80125 Napoli, Italy*

Correspondence should be addressed to Roger P. Karrer, roger.karrer@credit-suisse.com

Received 25 January 2008; Accepted 7 August 2008

Recommended by Jong Hyuk Park

Wireless mesh networks have the potential to provide ubiquitous high-speed Internet access at low costs. The good news is that initial deployments of WiFi meshes show the feasibility of providing ubiquitous Internet connectivity. However, their performance is far below the necessary and achievable limit. Moreover, users' subscription in the existing meshes is dismal even though the technical challenges to get connectivity are low. This paper provides an overview of the current status of mesh networks' deployment, and highlights the technical, economical, and social challenges that need to be addressed in the next years. As a proof-of-principle study, we discuss the above-mentioned challenges with reference to three real networks: (i) *MagNets*, an operator-driven planned two-tier mesh network; (ii) *Berlin Freifunk* network as a pure community-driven single-tier network; (iii) *Weimar Freifunk* network, also a community-driven but two-tier network.

1. INTRODUCTION

Wireless networks have the potential to realize the long-standing vision of ubiquitous high-speed Internet access. Therefore, they may revolutionize society in the 21st century, as the transistor and the Internet did in the 20th century, since the ubiquitous availability of information and communication will change the way we communicate with people and machines. Moreover, wireless technologies will also foster the availability of Internet services in rural areas and close the digital divide.

Today, we are in the middle of the deployment of wireless mesh infrastructures, and therefore also in the middle between initial hype and real numbers in terms of technical and economic feasibilities. Thus, we believe that this is the perfect time to take a step back and look at the current status of wireless mesh networks (WMNs) [1, 2].

In the first part of this paper, we assess whether the hype of realizing a ubiquitous high-speed Internet access is being realized, or whether reality is biting back. Can the technical specifications and algorithms live up to the expectations and visions? Are users jumping on the great features of mesh networks as predicted? To anticipate some of our findings, we will show that the first generation of mesh networks that are being deployed in cities shows the feasibility of wireless mesh networks to provide ubiquitous access. However, unfortunately, the performance of the networks is dismal; experience shows that the throughput is limited, and unfairness and throughput degradations of multihop communication impose severe limitations [3]. Moreover, from an economical perspective, subscription rates to city-wide meshes, such as in San Francisco, are dismal. Even though the fees are just a few dollars per month for a flat rate access of several Mbps, the subscriptions are far below the expectations.

In the second part of the paper, we leverage our findings about the current status to derive the challenges for what we call second generation of mesh networks. At a technical level, we must find means to scale the throughput to Gbps by a combination of hardware improvements as well as specialized algorithms for mesh networks. At an economical level, wireless mesh networks must find a feasible position between the established and extreme positions that we find today: wired networks with their high bandwidth and predictable performance on one side and 3G networks with their nation-wide coverage. Will wireless mesh networks continue to run in unlicensed spectrum or is it necessary to allocate licensed spectrum for meshes?

Finally, in the last part, we reflect the status and the challenges in three case studies. In particular, we discuss the technical, economical, and social challenges and differences in the *MagNets*, an operator-driven planned two-tier mesh

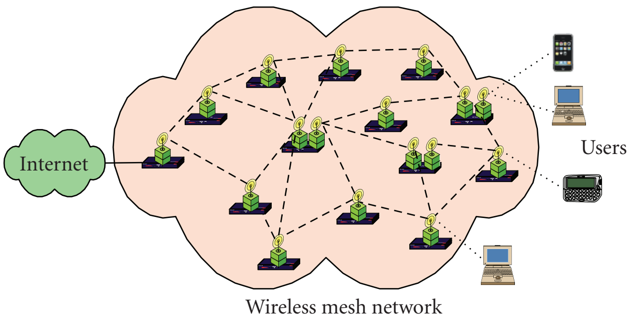

Figure 1: Mesh network.

network, the *Berlin Freifunk* network as a pure community-driven single-tier network, as well as in the *Weimar Freifunk* network, also a community-driven but two-tier network.

Our conclusions are intentionally controversial to stimulate a discussion among researchers and industry. We argue that wireless mesh networks will not be deployed for user access—at least from an economic point of view. Instead, they will be financed to increase the automation of remotely controlled devices, such as meters for gas or heating, parking meters, and traffic lights, whereas the financial contributions of users will be dismal.

The remainder of this paper is organized as follows. Section 2 gives an overview of the current status of wireless mesh networks. Section 3 outlines the challenges that need to be addressed in the next years. Then, Section 4 presents the case studies of the three deployed mesh networks. Finally, we draw our conclusions in Section 5.

2. CURRENT STATUS OF WIRELESS MESH NETWORKS

The wireless mesh networks we consider in this paper can be defined as an aggregation of infrastructure-based, wire-powered, stationary nodes that are equipped with at least one wireless card, as depicted in Figure 1. Some nodes, but not all of them, are additionally equipped with a wired Internet connection (e.g., DSL). The aggregation of nodes collaborates to provide coverage to an entire area, such as a university campus or an entire city, by forwarding data from a user who is attached to any of the mesh nodes over multiple wireless hops towards one of the mesh nodes that has a wired Internet link. Thus, we can divide the functionality of the nodes into two parts: to provide connectivity to users attached to the node, and to forward data from and to the wired mesh nodes. The latter is often termed as the "backhaul" of a wireless mesh network.

Compared to other definitions of mesh networks, we deliberately exclude the idea that user terminals (e.g., laptops) can be used to even further extend the coverage of the mesh by forwarding data from another user to an access point. Even though such an extension is technically possible, we exclude it for three reasons. First, laptops must be configured accordingly to forward the data. This configuration is beyond the control of the infrastructure mesh; instead it must be configured by users. Second, it is unlikely that users will dedicate their resources especially battery, but also CPU and network resources, to others unless they receive some benefit. Instead, such an operation incurs security risks. Third, users may turn on and off their laptops at any time, or also move around. Taking mobility and frequent topology changes into account increases the complexity of the mesh without the promise of significant performance gains.

Today, we see a plethora of mesh networks being deployed for research purposes but also as production networks in cities. After the seminal work by the MIT Roofnet [4], a large number of universities provide campus coverage via mesh networks. Next, efforts by Rice University have fostered the Technology-for-All (TfA) network in Houston, Tex, USA, which provides connectivity to underprivileged neighborhoods, with the vision to reduce the digital divide [5]. Finally, lots of cities worldwide plan or have deployed a city-wide WiFi mesh, including San Francisco, Singapore, London (the center, mostly for business customers), or Venice (for tourists).

Does this wave (or even flood) of deployment imply that wireless mesh networks have addressed all their challenges? These only minor questions in research and productive deployment are left! Quite interestingly, we find quite the opposite; namely, the current mesh networks are far from achieving sufficient quality in terms of performance and reliability, the security is in its infancy, and the economical aspects of wireless mesh networks raise more questions after the initial deployments than before. The remainder of this section discusses these issues in detail. In particular, we also take the survey by Akyildiz et al. [1] as a reference, and point out the differences and advances over the last 3 years.

2.1. *Quality*

The critical design factors that determine the quality of a wireless mesh network are performance, reliability, and scalability. Performance starts at the physical layer where the hardware defines the maximal capacity of a link. Current state-of-the-art WiFi cards and access points achieve a net throughput of 54 Mbps, as defined by the 802.11a/g standards. Capacity enhancements have been promised with 802.11n, where directional and smart antennas as well as MIMO and multiradio/multichannel systems promise rates of up to 600 Mbps.

Thus, it seems that at least the lower layers are on a good path towards the envisioned Gbps speeds. But how much of this capacity is available at the application level? The protocol overhead of the current Internet stack accumulates for roughly 50% of the capacity, implying that an approximate of 30 Mbps can be achieved. But are these the numbers we see in today's wireless networks? Fortunately, the MagNets outdoor network in Berlin shows link speeds of 30 Mbps on one link, over 500 m with directional antennas [6]. However, out of the 6 links in the testbed, only one link achieves this throughput because multiple conditions must be fulfilled to achieve this high throughput: perfect line of sight, directional antennas, and no interference. In fact, the link is based on 802.11a technology, and the number of interfering networks in the 5 GHz frequency band is still low. The other links in the MagNets testbed achieve between 16 and 18 Mbps. Unfortunately, the MagNets backbone is an exception in terms of performance, as many other deployed networks achieve only single-digit throughputs; for example, the TfA network has a throughput of 6–8 Mbps.

These throughputs are achieved with directional antennas and dedicated mesh nodes that form the backhaul of a mesh network. However, many mesh nodes available today at reasonable costs are equipped with a single WiFi card. This WiFi card must then be shared for 2 purposes: to forward data along the backhaul, and to service the users attached to the node. Since each operation requires both the receiving and the sending of data and only one operation is possible concurrently, the measured throughput of WiFi meshes that rely just on a single WiFi card are often limited to 1-2 Mbps.

Apart from poor performance, mesh networks suffer from multihop performance degradation and unfairness [7]. Multihop performance degradation, that is, the fact that traffic that is forwarded over multiple hops receives only a fraction of the throughput that a single-hop flow achieves, occurs because of the random access of the MAC protocol. A flow that traverses multiple hops has to compete multiple times in order for the medium to reach the destination. With existing 802.11 protocols, each competition is fair, such that the probability that a multihop flow packet reaches the destination is significantly lower than that of a single-hop flow. This issue is well known, and is expected to be addressed in the upcoming 802.11s standard for mesh networks.

Going up one layer in the hierarchy, routing in mesh networks is still an active area of research. Over the past decade, a plethora of routing protocols has been proposed for ad hoc networks. However, these protocols are conservative, pessimistic, and simplistic in their behavior because they consider that nodes may come and leave. In contrast, for mesh networks that are infrastructure-based, routing protocols are needed, which scale to larger areas and to a larger number of flows and rely on different metrics. Most ad hoc routing protocols rely on hop count as a metric. However, this metric is not suited for all applications, and does not guarantee the best usage of the underlying capacity.

At the transport layer, mesh networks can incur severe performance degradations, particularly as a function of the underlying routing protocol. Current implementations of TCP are prone to packet reordering, and react to variations in the delay. Thus, from a TCP point of view, all lower-layer protocols should try to conserve the routes (e.g., via static routing). Thus, these demands are exactly the opposite requirements of the network layer, where packets should be forwarded as dynamically as possible over different routes to opportunistically exploit channel fluctuations.

In summary, we realize that in fact most questions related to wireless mesh networks are largely unaddressed. In particular, when we require that answers to the above questions be not only written down as paperware but be evaluated in wireless mesh testbeds, we realize that we are worlds away even from understanding the behavior of wireless mesh networks—let alone being able to run them efficiently.

2.2. Security

Security in mesh networks still lacks efficient and scalable solutions. This dark observation stems in part from the fact that the Internet architecture lacks built-in security mechanisms. Thus, wireless mesh networks “inherit” the security properties/drawbacks of the Internet, and are therefore prone to flooding, DDoS attacks, and other malicious operations. In addition, however, wireless mesh networks add the drawbacks of the underlying wireless medium. Jamming attacks that prevent data transmissions from any wireless node in the neighborhood as well as attacks that exploit the features of the MAC, such as backoff procedures and network allocation vector (NAV) value settings in addition to blackhole routing where the attackers advocate routes to neighboring mesh nodes but just discard all received packets, are just examples of attacks that are easily mounted in wireless environments. As an addition to the negative tunes, the approaches known from the wired world, such as adding AAA (i.e., authentication, authorization, and accounting), are ill-suited for mesh networks because there is, and should be, no central service in a mesh work.

2.3. Economy

One of the key advantages of mesh networks has always been the low deployment costs [8]. While these arguments still hold today, we have learned over the last few months that they are not sufficient. In particular, on one hand, wireless mesh networks combine the advantages of the speeds of wired networks with the coverage of cellular networks. However, if we look at wireless mesh networks from a customer-consumer perspective, these advantages seem to turn into disadvantages. If a user is to pay for access, it is likely that the user chooses a fixed line at home and a cellular phone where connectivity is available worldwide. From this perspective, it seems that wireless mesh networks do not offer sufficient advantages to either justify yet another expense for connectivity or to even replace one of the other connections with WiFi.

These experiences are reflected in the news from San Francisco. In Spring 2007, EarthLink, the provider that runs the San Francisco network, reported a 30-million-dollar loss and a dismal subscription of 2000 users only. Moreover, the users and authorities are increasingly growing aware of privacy issues for the users, as Earthlink and Google may collect information about the location of the users and the sites they visit [9].

3. CHALLENGES

Based on the above analysis, we identify significant shortcomings in currently deployed wireless mesh networks. We believe that these deficiencies have only occurred in the first generation of wireless mesh networks that focused on providing the proof of concept for wireless mesh networks. However, these deficiencies must be addressed in the second generation of wireless networks. The remainder of this section highlights the challenges, and points out possible solutions.

3.1. Quality

The quest to achieve performance, reliability, and scalability in wireless mesh networks must be concurrently started at all layers. At the physical layer, improvements are on their way with multiple antenna systems, orthogonal frequency-division multiplexing (OFDM), and with novel 802.11 flavors such as 802.11n. In addition, however, two alternative research paths must be pursued. One is new wideband transmission schemes beyond OFDM and UWB (ultra-wide-band). These schemes must achieve higher transmission rates, and therefore push the capacity limits. Second, enhanced power schemes are needed to address the increasing interference. With the rapid deployment of wireless technologies in homes and cities, the degree of interference is constantly mounting. In the city of Berlin, during our measurements with the MagNets testbed [10], we have found up to 25 interfering networks in the neighborhood of one access point—per channel! Moreover, we have learned during the past two years that interference is the main reason for performance degradations, and not multipath fading. Thus, it is vital that interference is reduced by flexibly adjusting the power of wireless senders.

Tightly coupled with the physical-layer needs, there are the set of demands at the MAC layer. While advances at the physical layer provide the basic mechanisms, the MAC layer must determine how to use these mechanisms. For example, under which conditions the power should be increased or decreased to tradeoff the probability of correct reception of one packet against the interference with other neighboring access points. A strategy where everybody keeps the transmission power to its maximum is simply not going to work. Therefore, an enhanced collaboration between physical and MAC layers is required. A second set of work must deal with innovative MAC protocols. The current random access protocol, such as carrier sensing multiple access/collision avoidance (CSMA/CA), is far from being efficient and fair. Is a time division multiple access (TDMA) approach better, and in particular is it feasible when the schedule must take multiple distributed nodes into account? On the other hand, a TDMA solution would solve many issues. In particular, for ISPs, a TDMA solution would allow them to offer service-level agreements and have different service classes. These guarantees are necessary to create the desired revenues from mesh networks. Moreover, TDMA systems are likely to allow for a simple solution to the multihop unfairness and performance degradations.

At the network layer, the key challenge is to optimize the usage of the underlying capacity. This task is extremely challenging given the need to coordinate multiple distributed mesh nodes and given the wide heterogeneity of underlying mesh nodes and channels. What kind of routing metrics does show the best performance and best match the application needs? Is multipath routing a way to optimize the capacity usage? How can we integrate routing in a mesh with routing in the Internet? All these questions require a fundamental analysis and experimental evaluation before they can be answered. However, we note a recent interest in multipath routing or, to formulate it in a more general way, in diversity. Even in the Internet, the concept that only a single path is used through the Internet is currently questioned because it is likely that alternative paths exist, which may be less loaded and therefore have a better application-level performance. If the concept of diversity was integrated as a fundamental concept into a future Internet architecture, it could also help to improve the performance in a wireless mesh network.

At the transport layer, we face two challenges. At the actual stage, we know that current TCP implementations do not perform well over multihop wireless networks. Thus, it is necessary to tune and adapt TCP mechanisms to deal with large round trip time (RTT) variations, path asymmetries, and varying channel conditions at different time scales. The challenge thereby is to come up with solutions that achieve a high throughput in both wired and wireless networks, or to have different TCP implementations and find a way to dynamically choose a specific implementation based on the underlying network.

Finally, at the application layer, we see one dominant question, that is, whether there is such a thing as a killer application for mesh networks. It is unlikely that current applications require significant changes in their behavior depending on whether they are deployed over a mesh network or a wired network. It can be assumed that the lower-layer protocols take care of the difference. That is, VoIP applications require a routing based on delay minimization, whereas multimedia applications or peer-to-peer applications are likely to prefer routing protocols that achieve a high bandwidth. However, a killer application would push the limits and the requirements of future mesh networks into a specific direction.

Towards achieving the above goals, we should be aware that three types of work are required to make progress. First, at a theoretical level, work is required to help us understand the behavior of protocols. For example, we still ignore to a large degree how 802.11 MACs perform over multihop backhaul networks in real networks. That is, how exactly is data forwarded from one hop to another? This knowledge is vital to, for example, foster new MAC-layer protocols that rely on random access but do not have severe throughput and unfairness drawbacks. Second, novel protocols are needed that *significantly* improve the performance. In research, we often see research proposals that achieve 10 or 20% of improvements. Such small advances do not help us make progress. Instead protocols are needed, which double, triple, or n-ple the throughput. Finally, we need solutions that are experimentally evaluated and tested under several conditions. Over the last decades, for example, a plethora of routing protocols or enhancements thereof has been proposed. However, we still ignore how they would perform in a real network. In fact, they often perform well under a specific constraint but have severe drawbacks under others. It is vital for the progress that protocols are experimentally evaluated.

3.2. Security

Providing security must be one of the most dominant objectives in wireless mesh network research in the near future. Without securing wireless networks properly, it is likely that users will not use wireless mesh networks, as seen in the case of San Francisco. But how to secure a wireless mesh network? The good news is that security in wireless mesh networks often coincides with security in wired networks. Because the topology is known, mesh nodes know their neighbors and can ask for identification. Currently, the worst attack scenario is probably jamming, as jamming (all frequencies) does not leave room for automated solutions. However, the advantage is that jamming networks require that the attacker be near the mesh or that a jamming device be installed near the mesh. In either case, the jamming device can easily be identified by following the radiation pattern.

For all other attacks, we repeat the requirements by Yang et al. [11]. In future work, the main directions are as follows: (i) to critically evaluate any proposed security solution, including vulnerability analysis and measurements and emulations, and (ii) security protocols must be resilient and robust, possibly even against unknown attacks. By no means must a security protocol proposal make idealistic assumptions.

3.3. Economy

At an economical level, we identify three key directions. First, protocols and mechanisms must be implemented into wireless mesh networks to provide *carrier-grade services*. These services are a vital requirement for ISPs to create revenues. To enable carrier-grade services, protocols must be designed to achieve a predictable performance and allow for quality differentiation. At the MAC layer, TDMA could be an option, but similar efforts are required at all levels. For example, streaming services must be deployed. Moreover, AAA and related mechanisms must be built into meshes. In contrast to wired networks where service guarantees are achieved with overprovisioning today, it is clear that such an approach is not feasible in a wireless world—at least not by scaling bandwidth.

Second, related to carrier-grade services is the following question. How much frequency is needed for wireless technology? As discussed above, the increasing deployment of wireless technology incurs interference and is therefore already now the main “killer” of performance. Adding more spectrum certainly helps. The key question thereby is as follows. Should the spectrum continue to be free, or should it be licensed? Clearly for a TDMA system to work, a licensed spectrum is a precondition, as otherwise any random access technology in the same frequency band would interfere with the TDMA schedule. Discussions about issuing small frequency bandwidth to ISPs for a relatively low cost are already ongoing in different countries.

Third, the killer application for meshes must be found. Actually, there are two types of killer applications: the killer application that motivates the deployment of mesh networks, and the killer application for users to use the mesh. These two applications may be different or can be the same. For the killer application that motivates the deployment, the use of this application must create revenues or savings that compensate for the investment of mesh deployment. Potential killers here are the meters for gas, heating, power or parking, and remote surveillance and emergency situations. For example, if all meters were equipped with cheap WiFi senders, their level could remotely be controlled, saving the costs of sending people to homes. Remote surveillance and emergency may help police, fire departments, and ambulances to get a picture of an emergency situation at an early stage and prepare the rescue accordingly. For users, video and TV streaming is often considered as the killer application. However, are we really all such addicted to TV that we need to receive streams at high data rates all the time? Or do location-based services find the right balance between providing useful information and ensuring the privacy of users? Thinking along these lines, it seems that the technological challenges are far better understood than the demands of the users and the society.

4. CASE STUDIES

This section describes and studies the status and the challenges of three deployed wireless mesh networks:

(i) the operator-driven and planned *MagNets* network in Berlin;

(ii) the community-driven one-tier *Berlin Freifunk* network;

(iii) the community-driven two-tier *Weimar Freifunk* network.

The analysis provided in this section aims at showing the wide variety of *technical*, *economical*, and *social* motivations, parameters, and goals behind mesh networks, and therefore also reveals the tradeoffs among them. The text describes the meshes in detail, and Table 1 gives an overview of the comparison. More precisely, in this section—for each considered mesh—we first report a description of the network and then we provide details, respectively, on quality, security, and economics. This analysis highlights the main differences between operator-driven and community-driven networks.

4.1. MagNets

The *MagNets* project aims at deploying a *semiproductive testbed*, that is, a testbed where we perform experimental research of protocol behavior but where at the same time university students use the network as an operational one to get access to the Internet (*MagNets* is a short form of Magenta networks, where Magenta is the trademark color of Deutsche Telekom; see http://www.deutsche-telekom-laboratories.de/~karrer/magnets.html). The objective of *MagNets* is to get as close as possible to the vision of high-speed wireless ubiquitous Internet access, with carrier-grade quality and support of service-level agreements. For this

Table 1: Comparison of wireless mesh networks in Germany.

Evaluation	Parameter	*MagNets*	*Freifunk Berlin*	*Freifunk Weimar*
Technical	Deployment	Planned	Community	Planned community
	Architecture	2-tier	1-tier	2-tier
	Mode	802.11super a/g	802.11bg	802.11abg
	Antennas	Omni-/directional	Omni-	Omni-
	Backhaul	Directional, freq. separated	N.A.	Freq. separated
	Bandwidth	<62 Mbps	<13 Mbps	<15 Mbps
	Number of nodes	100	800+	150
	Number of gateways	Variable (up to 80)	15	15
	Gateway line speed	100 Mbps	1–16 Mbps	1–6 Mbps
	Average nodes per gateway	1–5	55	15
Economics	Deployment	Planned	Community	Planned community
	Cost per node	>300 $	50$	50$
Social	Line lease	University	Users	Community
	AP authentication	List	Free	List/trust
	User authentication	802.1x	Free	List/trust
	Firmware distribution	Central	Distributed	Central
	Firmware status	Homogeneous	Heterogeneous	Homogeneous
	Services	Internet, IPTV, etc.	Internet	Internet

purpose, *MagNets* is designed as a two-tier architecture, with a designated high-speed wireless backbone and an access tier. While the access tier supports standard 802.11 with omnidirectional antennas, we focus in particular on the high-speed backbone that shows interesting and distinguishing features.

The backbone consists of 5 nodes that connect high-rise buildings in the heart of Berlin over a total distance of 2.3 km, as depicted in Figure 2. Each node consists of a Linux router and one access point per outgoing link that is connected to a directional antenna. Therefore, data can concurrently be sent over all links, and the directional antennas reduce the interference and therefore allow for spatial reuse. Two access points support 802.11 SuperAG mode, supporting up to 108 Mbps. Two links operate in the 5 GHz range, while the others operate in the 2.4 GHz range. The transmission and throughput capability of a *MagNets* backbone node significantly exceeds that of a “traditional” mesh node that consists of a single access point with a single WiFi card. More information on the backbone can be found in [6, 10, 12].

4.1.1. Quality

Figures 3 and 4 show the throughput of links 1 (5 GHz) and 3 (2.4 GHz), respectively. In basic mode (802.11ag), the application-layer throughput is 31 Mbps for link 1, and 8.4 Mbps for link 3. Given that the raw throughput is 54 Mbps in the basic mode and that 50% have to be deducted for protocol and messaging overhead, link 1 is close to the optimal performance. In contrast, the performance of link 3 is significantly due to interference of competing networks. However, by putting the nodes into SuperAG mode, which results in 108 Mbps raw throughput, we note that the throughput on link 1 achieves 62.4 Mbps and 50.3 Mbps on link 3. Detailed results and discussions on achievable performance of *MagNets* can be found in [6, 12].

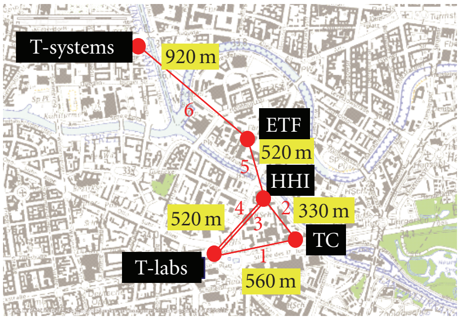

Figure 2: *MagNets* WiFi backbone in the heart of Berlin.

Among the applications that can be supported by such a backbone, there is IPTV. In particular, we were addressing the problem that many users may want to watch TV on their mobile devices, such as laptops or iPods. These devices are equipped with WiFi, but not with other interfaces that allow the reception of TV. On the other hand, in Berlin, DVB-T is available throughout the city and can be received with USB receivers. Our idea was thus to use the mesh as a technical relay by placing one DVB-T receiver into the mesh, converting the DVB-T signal into IP packets, and distributing the TV stream to the users via WiFi [13]. Figure 5 shows the frame rate of correctly received frames at a client connected to the backbone as a function of time. The figure shows that the backbone is able to maintain an almost reliable frame rate. The average frame rate is 28 frames per second, out of 30 transmitted, with a standard deviation of 2

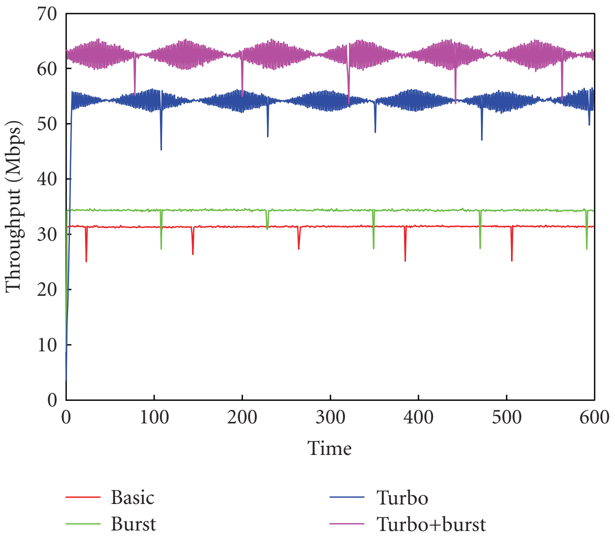

Figure 3: Throughput on link 1 (5 GHz).

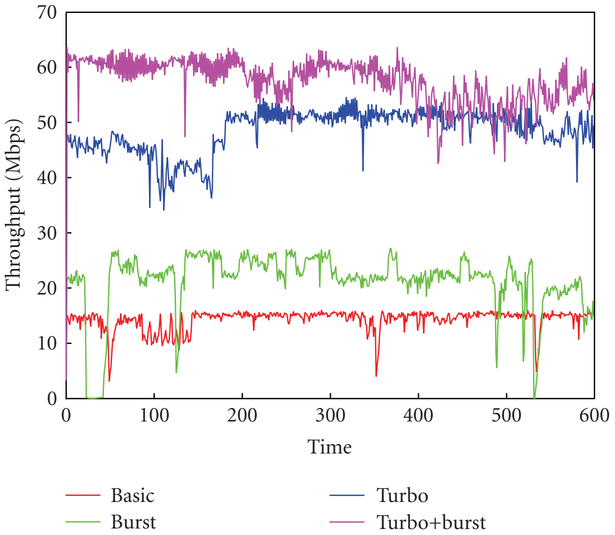

Figure 4: Throughput on link 3 (2.4 GHz).

frames per second. These rates clearly lead to an acceptable if not excellent viewing experience by a user.

Thus, we note that the planned deployment and the high-power hardware per node result in high per-link and multihop throughput. More specifically, the measured application-layer throughput is close to the optimal achievable throughput, and the link quality is high throughout the entire measurement time. Therefore, the backbone is able to provide high-speed wireless Internet access.

4.1.2. Security

All nodes that are deployed as part of the *MagNets* network are managed by a single operator. The location of all nodes is stationary and well known. Most nodes are equipped with a (low-bandwidth) wired connection that is used for management purposes only, such as remote upgrades (since the nodes are wire-powered, it is typically easy to combine the power line with a cable for connectivity). Thus, since the nodes and their location are known, suited authentication schemes can be used to eliminate malicious nodes from being introduced into the mesh. Similarly, firmware upgrades and software installations are made over the wired management network. This leads to a significant reduction of the threat potential compared to, for example, community networks, as discussed below. However, it does not prevent adversaries from jamming attacks at the physical layer or DDoS attacks at the higher layers.

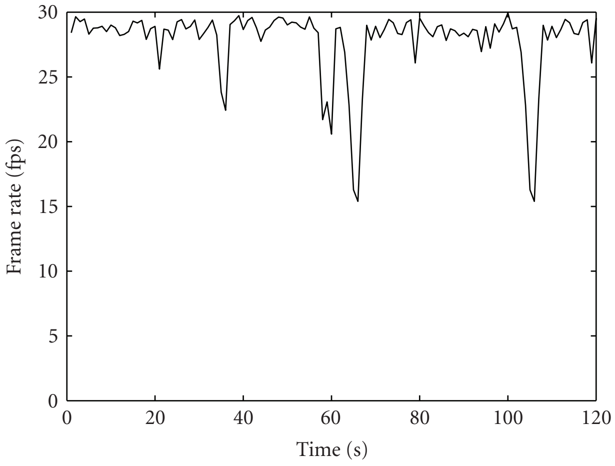

Figure 5: TV streaming over the backbone; frame rate at the client.

4.1.3. Economics

The deployment of such a high-speed WiFi network is costly. The costs per node are easily one order of magnitude higher than those of community networks. In concrete numbers, a *MagNets* backbone node is in the order of several 100 dollars, whereas the simple nodes used in the *Berlin Freifunk* are typically available for 50 dollars. Multiply the per-node costs by the number of nodes and add the deployment efforts, then the numbers begin to increase. It does therefore come to no surprise that operators are basically interested in WiFi meshes, but that the calculation of the deployment costs as well as the operational costs must be compared against the potential revenue. While it would be interesting to know the business cases for WiFi meshes, for example, the break-even point or the maximal cost per access point that would allow an operator to create revenue, these numbers are unfortunately not disclosed to the public. Given the bad news from the deployed mesh networks, we can only speculate that the costs are currently too high even though no spectrum costs arise.

4.2. Freifunk community mesh network in Berlin

Several kilometers away from *MagNets* lies one of the largest and biggest community mesh networks in the world in terms of nodes deployed and area covered: the *Berlin Freifunk* network. The project was born in 2003 out of the need to provide connectivity to the households in the former East Berlin area. For historic reasons, East Berlin is equipped with a state-of-the-art fiber network, but this network is not able to offer ADSL like copper networks. Therefore, a community effort was started to cover the area of East Berlin with a WiFi mesh. A core group of 5 main programmers set out to build firmware and software, and today their practical considerations on building and running mesh networks are highly respected, for example, their contributions to the optimized link state routing (OLSR) protocol. Thus, the motivations and goals are significantly different from the ones of *MagNets*.

4.2.1. Quality

As of January 2008, the *Berlin Freifunk* has reached the size of 820 participating nodes. For a small part of the network topology, the density of the nodes and the links are depicted in Figure 6. Due to the organic growth of the network, the mesh structure is flat; all nodes are transmitting on channel 10 (2457 GHz) in ad hoc mixed mode based on 802.11b and 802.11g. Thus, compared to *MagNets* and to the *Weimar* network described below, the *Berlin Freifunk* lacks an efficient two-tier structure and does not make use of the available spectrum in the 5 GHz range to avoid the overloaded spectrum in the 2.4 GHz range. Moreover, also the structure of the single tier is not planned, as mesh nodes are put up by individuals who join the network. Therefore, areas with a high node density (and thus high interference) coexist with areas with sparse connectivity.

Data is forwarded among the mesh nodes from and to currently 15 Internet active gateways, with ADSL line speed from 1 up to 16 Mbps. Thus, on average, 55 nodes share the line speed of a single Internet gateway, with local deviations that increase the per-gateway node even higher. The hop-count values to an Internet gateway vary from 1 up to 18 hops, with an average value equal to 5 hops.

The achievable application-layer throughput between two nodes is 13 Mbps in the best case, when the transmission rate on both nodes is 54 Mbps (802.11g). But, many nodes are still using mixed 802.11bg, and they are therefore a severe performance killer for the end-to-end throughput. Thus, in terms of quality, and also compared to MagNets, the quality of an end-to-end connection heavily depends on the hop count. Finally, from a mesh perspective, there is a severe unfairness towards clients that are more hops away from the Internet gateway. Obviously, with these low data rates, real-time applications are not supported.

4.2.2. Security

The *Freifunk* mesh network is basically free to use for everybody who is within the range of the network. The network entirely lacks technical access restrictions and mechanisms to regulate access to the network. The data transmissions over the wireless transmissions are not encrypted. Thus, no single technical mechanism is implemented to exclude misbehaving nodes or users. The network could be named as an *autarc* and insecure wireless network. Because of the decentralized administration of the nodes, the open source *Freifunk* firmware is not maintained and updated so that many different releases with potential security problems are active in the network.

The administration of a node is done either via an "ssh" session or a web interface based on "https." Node owners may also use their laptops without having the OLSR daemon. This access is based on the MAC address of the laptop, and is bounded to the node to which the user may have access.

The only centralized service that is required in the network is the IP allocation scheme. This scheme is not automatically configured, but users register their nodes on a central wikipedia; for the *Freifunk* project, it is used to coordinate the IP addresses' allocation activity. The IP addressing scheme is based on a 10.0.0.0/8 network, and the numbering schema is correlated to the different districts of Berlin.

4.2.3. Economics

Within such a mesh, the single-node owner can be seen as a kind of a mini provider who invested about 100 Euros in some common wireless router and who pays the energy costs to operate his node. The range of the mesh network increases with new participants and their packet forwarding ability. Internet gateways are provided by individual users without charging any fees for that service. Finally, there is no legal form of a company or a registered club in place. The mesh network in Berlin is more or less a *voluntary* network without any contract between the users. To summarize, we have what follows.

(i) A big community, with the technical center called *C-Base*, doing weekly workshops on how to build an antenna, setting up hardware, and configuring devices.

(ii) There is no registered club in place.

(iii) The participation in the mesh network and the use of the Internet connection are free.

(iv) The energy consumption is around 8–10 watt per node.

(v) Costs amount to 150 Euros per month per node.

4.3. *Freifunk community mesh network in Weimar*

Another big community mesh network has been set up in the city of Weimar, located in the south of Germany. The project was born in the end of 2003 out of the same need as in Berlin: to provide Internet access to the households where no broadband access was available. A core group, made up of one main programmer and two administrators, did the work for building a manageable mesh network based on the firmware provided by the Berlin group. The motivations to build the mesh are similar but the goals are different (with respect to the ones of the community in Berlin).

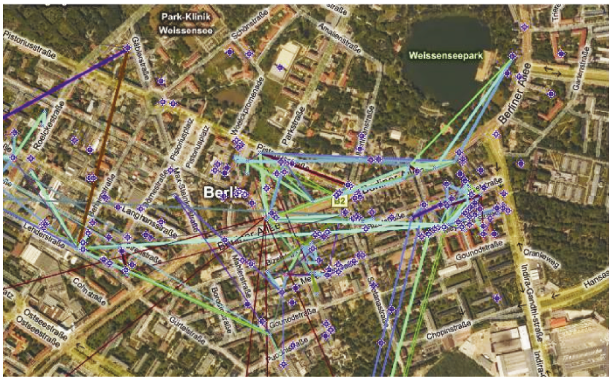

Figure 6: Map of *Berlin Freifunk* network.

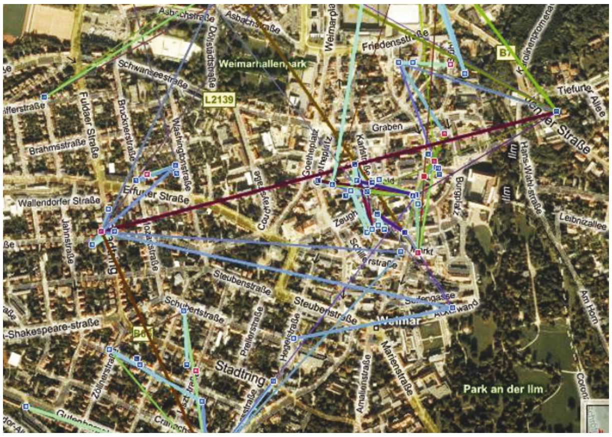

Figure 7: Map of *Weimar Freifunk* network.

4.3.1. Quality

As of January 2008, the Freifunk mesh in Weimar has reached the size of about 150 participating nodes. Figure 7 contains a cutout of the topology, and it shows the links and the node density. The topology of the network, when compared to the one in Berlin, is more structured. The mesh network consists of three main clouds with a high density of nodes running all on the same channel, channel 1 (2412 GHz) in 802.11g-only mode. Thus, compared to the Berlin network described above, the *Weimar Freifunk* does make use of the 5 GHz spectrum; the backbone, connecting the three main clouds, consists of five nodes wirelessly bridged together using 5 GHz. Another difference between *Freifunk Berlin* and *Weimar* is the better percentage of nodes per Internet gateway. Data is forwarded among the mesh nodes from and to 15 Internet active gateways, with ADSL line speed from 1 up to 6 Mbps. All these features permit the *Weimar Berlin* to reach higher performance when compared to the *Freifunk Berlin*.

4.3.2. Security

From the encryption point of view, both networks in Berlin and Weimar are similar because both do not use any encryption on the wireless interface at all. At the same time, the mesh network in Weimar presents some important differences. The network access is restricted to nodes which are registered on a central web page. This represents the basis of the so-called *white list* of allowed nodes. This technical mechanism is used to exclude misbehaving nodes or users. Furthermore, the firmware update is managed by a centralized process, which leads to a more homogeneous firmware distribution. As in Berlin, the administration of a node is done either via an “ssh” session or a web interface based on “https.” Node owners may also use their laptops without having the OLSR daemon. This access is based on the MAC address of the laptop, and it is not only bounded to the node to which the user may have access. More precisely, if the MAC address of a laptop is in the access list of at least three nodes of the mesh, this information is distributed via OLSR service announcements over all nodes, which leads to a mesh-wide access from such a laptop. In other words, three nodes are needed, trusting a certain MAC address of a laptop to provide a mesh-wide access to this laptop. The IP addressing scheme used in the mesh of Weimar is similar to the one in Berlin, based on a central user registration. Also, the IP addressing scheme is based on a 10.0.0.0/8 network, and the numbering schema is correlated to the different districts of Weimar.

4.3.3. Economics

From the cost structure point of view, the situation in Weimar is the same as in Berlin. Every user is responsible for purchasing and operating his own node. Besides the situation in Berlin, where voluntary users share their Internet connection with the community, in Weimar was founded a registered club called “Weimarnetz.” The task of this registered club is to rent the different ADSL Internet lines so that no single user is responsible for the activities on the network. As in Berlin, the use of the mesh network and, therefore, the use of the Internet connectivity are free of charge. The mesh network in Weimar is more or less a voluntary network without any contract between the users, but with some more centralized approaches to achieve a better network performance while having a more homogeneous network compared to the one in Berlin.

4.4. Discussion

In this paper, we have described and discussed the status and the challenges of three deployed wireless mesh networks: (i) the operator-driven and planned *MagNets* network in Berlin; (ii) the community-driven one-tier *Berlin Freifunk* network; (iii) the community-driven two-tier *Weimar Freifunk* network. The analysis provided in this work showed the broad assortment of *technical*, *economical,* and *social* motivations, features, and goals behind mesh networks, and therefore also revealed the tradeoffs and the main differences among them. In summary, Table 1 shows an overview of the different parameters and tradeoffs. Pure community-driven networks have a flat organization, lack security features, and achieve a low throughput, whereas operator-driven mesh networks aim at providing carrier-grade throughput, services, and security. In between these extremes, different options exist, such as the community-operated network in Weimar where a central organization manages the structure, connectivity, gateways, and access. Thus, we find that the motivations behind the mesh networks are different, and the resulting deployment and operation are therefore diverse as well.

5. CONCLUSIONS

This paper gave an overview of the current status of wireless technology and its deployment, in particular wireless mesh networks, as well as the challenges that are to be addressed in the near future. We considered three case studies: *MagNets*, *Berlin Freifunk* network, and *Weimar Freifunk* network.

Our findings show that current mesh networks show the feasibility of providing WiFi coverage to large areas, such as entire cities, but not much more.

First, at a technical level, current mesh networks are far from being efficient, and protocols at all levels must be developed to provide carrier-grade services that allow ISPs to create revenues from mesh networks and therefore compensate for the investments of the mesh infrastructure.

Second, as for the security, meshes are as much in the infancy as the wired world. However, without the protection of the wired medium, further protection is needed to ensure a secure data transmission. Finally, a key point in security is protecting the privacy of the users. The position of a user can easily be determined by the mesh node it connects to. It is far from being clear how and whether this privacy is sufficiently protected.

Third, at an economical level, mesh networks seem to combine the advantages of wired-like performance and cellular-like coverage. However, from a user's perspective who has to pay for connectivity, it rather looks as if mesh networks combine the disadvantages.

Thus, the stakes are high and the challenges are far from being easy to answer. Nevertheless, or exactly because of the challenges, we argue that wireless mesh networks still maintain a large research potential that is worth exploiting, mainly using experimental evaluations over real testbeds.

REFERENCES

[1] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless mesh networks: a survey," *Computer Networks*, vol. 47, no. 4, pp. 445–487, 2005.

[2] G. Bianchi, S. S. Chakraborty, X. Guo, E. Knightly, and D. Lee, "Multihop wireless mesh networks," *IEEE Journal on Selected Areas in Communications*, vol. 24, no. 11, pp. 1957–1958, 2006.

[3] 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 *Proceedings of the 22nd Annual Joint Conference on the IEEE Computer and Communications Societies (INFOCOM '03)*, vol. 3, pp. 1744–1753, San Francisco, Calif, USA, March-April 2003.

[4] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, "Architecture and evaluation of the mit roofnet mesh network," in *Proceedings of the 11th Annual International Conference on Mobile Computing and Networking (MobiCom '05)*, Cologne, Germany, August 2005.

[5] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier urban mesh access network," in *Proceedings of the 4th International Conference on Mobile Systems, Applications and Services (MobiSys '06)*, pp. 96–109, Uppsala, Sweden, June 2006.

[6] A. Botta, A. Pescapé, G. Ventre, and R. Karrer, "High-speed wireless backbones: measurements from magnets," in *Proceedings of the 4th International Conference on Broadband Communications, Networks and Systems (BROADNETS '07)*, pp. 680–689, Raleigh, NC, USA, September 2007.

[7] V. Gambiroza, B. Sadeghi, and E. W. Knightly, "End-to-end performance and fairness in multihop wireless backhaul networks," in *Proceedings of the Annual International Conference on Mobile Computing and Networking (MOBICOM '04)*, pp. 287–301, Philadelphia, Pa, USA, September-October 2004.

[8] R. Karrer, A. Sabharwal, and E. Knightly, "Enabling large-scale wireless broadband: the case for TAPs," in *Proceedings of the 2nd Workshop on Hot Topics in Networks (HotNets-II)*, Cambridge, Mass, USA, November 2003.

[9] A. Seybold, The big picture: where the industry is today and where it is headed.

[10] R. Karrer, I. Matyasovszki, A. Botta, and A. Pescapé, "MagNets—experiences from deploying a joint research-operational next-generation wireless access network testbed," in *Proceedings of the 3rd International Conference on Testbeds and Research Infrastructure for the Development of Networks and Communities (TridentCom '07)*, pp. 1–10, Orlando, Fla, USA, May 2007.

[11] H. Yang, F. Ricciato, S. Lu, and L. Zhang, "Securing a wireless world," *Proceedings of the IEEE*, vol. 94, no. 2, pp. 442–454, 2006.

[12] R. P. Karrer, I. Matyasovszki, A. Botta, and A. Pescapé, "Experimental evaluation and characterization of the magnets wireless backbone," in *Proceedings of the 1st ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH '06)*, pp. 26–33, Los Angeles, Calif, USA, September 2006.

[13] R. P. Karrer and D. Reim, "Ubiquitous TV delivery to the masses," in *Proceedings of the International Conference on Intelligent Pervasive Computing (IPC '07)*, pp. 155–160, Jeju, Korea, October 2007.