2nd generation wireless mesh networks: technical, economical and social challenges

Roger P. Karrer
Deutsche Telekom Laboratories
TU Berlin
D-10587 Berlin, Germany
roger.karrer@telekom.de

Antonio Pescape
University of Napoli
Federico II
Napoli, Italy
pescape@unina.it

Abstract— **Wireless mesh networks have the potential to provide ubiquitous high-speed Internet at low costs. The good news is that initial deployments of WiFi meshes show the feasibility of providing ubiquitous Internet access. 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 network deployment and highlights the technical, economical and social challenges that need to be addressed in the next years.**

I. 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, as 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 feasibility. Thus, we believe that this is the perfect time to take a step back and look at the current status of WMNs (Wireless Mesh Networks) [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 show 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 multi-hop communication imposes severe limitations [6]. Moreover, from an economical perspective, subscriptions rate 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?

Our conclusions are intentionally controversial to stimulate a discussion among researchers and industry. We argue that wireless mesh deployments 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, traffic lights, etc, whereas the financial contributions of users will be dismal.

The remainder of this paper is organized as follows. Section II gives an overview of the current status of wireless mesh networks. Section III outlines the challenges that need to be addressed in the next years. Finally, we draw our conclusions in Section IV.

II. 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 that 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

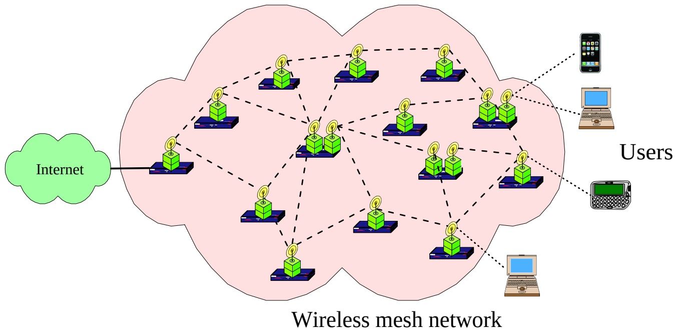

Fig. 1. Mesh Network.

to forward data from and to the wired mesh nodes. The latter is often termed 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, for others, unless they receive some benefit. Instead, such an operation incurs security risks. Third, users may turn on and off their laptop 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 [3], 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 TX that provides connectivity to underprivileged neighborhoods, with the vision to reduce the Digital Divide [5]. Finally, lots of cities world wide 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? That only minor questions in research and productive deployment are left? Quite interestingly, we find quite the opposite: namely that current mesh networks are far from achieving sufficient quality in terms of performance and reliability, that security is in its infancy, and that 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.

A. Quality

The critical design factors that determine the quality of a wireless mesh network are performance, reliability and scalability. Performance start 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, MIMO and multi-radio/multi-channel 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 $500m$ with directional antennas [4]. 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, e.g. 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: forward data along the backhaul and 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 $1-2$ Mbit/s.

Apart from poor performance, mesh networks suffer from multi-hop performance degradation and unfairness [7]. Multi-hop performance degradation, i.e. 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 for the medium to reach the destination. With existing 802.11 protocols, each competition is fair, such that the probability that a multi-hop 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 that scale to larger areas and to a larger number of flows, and that 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, in particular as a function of the underlying routing protocol. Current implementations of TCP are prone to packet reordering and it reacts to variations in the delay. Thus, from a TCP's 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 be able to run them efficiently.

B. *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, attacks that exploit the features of the MAC, such as backoff procedures and network allocation vector (NAV) value settings, 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. Adding to the negative tunes is that approaches known from the wired world, such as adding AAA (Authentication, Authorization and Accounting) are ill suited for mesh networks because their is, and should be, no central service in a mesh work.

C. *Economy*

One of the key advantages of mesh networks has always been the low deployment costs [9]. While these arguments still hold today, we have learned over the last few months they they are not sufficient. In particular, on the 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's and consumer's 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 world-wide. 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 [10].

III. 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.

A. *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, OFDM (Orthogonal Frequency-Division Multiplexing) and with novel 802.11 flavors, such as 802.11n. In addition, however, two alternative research paths must be pursued. One is new wide-band transmission schemes beyond OFDM and UWB (Ultra-wideband). 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 [8], we have found up to 25 interfering networks in the neighborhood of one access point - per channel! Moreover, we have learned in the past 2 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 trade off 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 layer is required. A second set of work must deal with innovative MAC protocols. The current random access protocol, such as CSMA/CA (Carrier Sensing Multiple Access/Collision Avoidance), is far from efficient and fair. Is a TDMA (Time Division Multiple Access) 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 multi-hop unfairness and performance degradation.

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 show the best performance and best match the application needs? Is multi-path 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 multi-path 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 that may be less loaded and therefore have a better application-level performance. If the concept of diversity were 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 multi-hop wireless networks. Thus, it is necessary to tune and adapt TCP mechanisms to deal with large RTT (Round Trip Time) 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: is there such a thing as a killer application for mesh networks. It is unlikely that current applications require significant changes in their behavior depending on if 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 that help us to understand the behavior of protocols. For example, we still ignore to a large degree how 802.11 MACs perform over multi-hop backhaul networks in real networks. I.e. how exactly is data forwarded from one hop to another? This knowledge is vital to e.g. 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 to make progress. Instead, protocols are needed that double, triple or n-ple the throughput. Finally, we need solutions that are experimentally evaluated and tested under several conditions. Over the last decades, e.g., a plethora of routing protocols or enhancements thereof have 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 in others. It is vital for the progress that protocols are experimentally evaluated.

B. *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 requires that the attacker is near the mesh or that a jamming device is 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 (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.

C. *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 that 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. E.g. 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, and related to carrier-grade services, is the 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: 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 application 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 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.

IV. Conclusions

This paper gave an overview of the current status of wireless technology and their deployment, in particular wireless mesh networks, and the challenges that are to be addressed in the near future. Our findings show that current mesh networks show the feasibility of providing WiFi coverage to large areas, such as entire cities. But not more. At a technical level, current mesh networks are far from efficient, and protocols at all levels must be developed that provide carrier-grade services that allow ISPs to create revenues from mesh networks and therefore compensate for the investments of the mesh infrastructure. Security-wise, 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 clear how and whether this privacy is sufficiently protected. Finally, 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 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. Important is that the exploitation is not incremental paperware, but is driven by visions that result in a giant leap ahead in our knowledge and that the results are analyzed and verified with experimental evaluations over real testbeds.

References

[1] I. Akyildiz, X. Wang, and W. Wang. Wireless mesh networks: A survey. *Computer Networks Journal (Elsevier)*, 47:445–487, Mar. 2005.

[2] G. Bianchi, S. Chakraborty, X. Guo, and E. Knightly. Multihop wireless mesh networks. In *IEEE Journal on Selected Areas in Communications*, Vol. 24, N. 11, Nov. 2006.

[3] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of the mit roofnet mesh network. In *ACM Mobicom'05*, Cologne, Germany, Aug. 2005.

[4] A. Botta, A. Pescapé, G. Ventre, and R. Karrer. High-speed wireless backbones: measurements from magnets. In *Proceedings of IEEE Broadnets'07*, Raleigh, NC, Sept. 2007.

[5] 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.

[6] 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.

[7] V. Gambiroza, B. Sadeghi, and E. Knightly. End-to-end performance and fairness in multihop wireless backhaul networks. In *Proceedings of IEEE MobiCom*, Philadelphia, Sept. 2004.

[8] 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 IEEE Tridentcom'07*, Orlando, FL, May 2007.

[9] R. Karrer, A. Sabharwal, and E. Knightly. Enabling large-scale wireless broadband: the case for TAPs. In *Proceedings of HotNets-II*, Cambridge, MA, Nov. 2003.

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

[11] H. Yang, F. Ricciato, S. Lu, and L. Zhang. Securing a wireless world. *Proceedings of the IEEE*, 92(2):442–454, Feb. 2006.