ACCEPTED FROM OPEN CALL

IEEE 802.11s: The WLAN Mesh Standard

Guido R. Hiertz, RWTH Aachen University
Dee Denteneer, Philips Research
Sebastian Max, RWTH Aachen University
Rakesh Taori, Samsung Electronics Co. Ltd.
Javier Cardona, cozybit Inc.
Lars Berlemann, T-Mobile International
Bernhard Walke, RWTH Aachen University

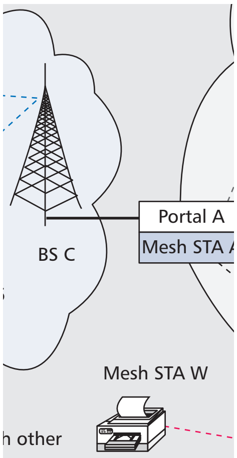

The authors provide insights into the latest developments in 802.11s and explain how the overall mesh concept fits into the 802 set of networking standards.

Abstract

The wireless local area network standard IEEE 802.11 is the preferred solution for low-cost data services. Key to its success are the 2.4 and 5 GHz unlicensed bands. The transmit power limitations imposed due to regulatory requirements limit the range (coverage) that can be achieved by WLANs in these bands. However, the demand for "larger" wireless infrastructure is emerging, ranging from office/university campuses to city-wide deployments. To overcome the limitations of single-hop communication, data packets need to traverse over multiple wireless hops, and wireless mesh networks are called for. Since 2004 Task Group S has been developing an amendment to the 802.11 standard to exactly address the aforementioned need for multihop communication. Besides introducing wireless frame forwarding and routing capabilities at the MAC layer, the 802.11s amendment brings new interworking and security. In this article, we provide insights into the latest developments in 802.11s and explain how the overall mesh concept fits into the 802 set of networking standards.

Introduction

Wireless LANs (WLANs) are proliferating and the desire for ubiquitous wireless connectivity is driving the demand for coverage extension of today's WLANs. However, regulatory limitations restrict the transmission power of WLAN devices.

We have been here before. *Bridging* evolved the Ethernet (802.3) standard from a single-hop to a multihop system. With bridges, the communication between end stations is no longer limited to the same LAN. WLANs are in an early stage compared to their wired ancestors. The present 802.11 interconnections rely on wired networks to carry out bridging functions. For a number of reasons, this dependency on wired infrastructure must be eliminated. First, this dependency is costly and inflexible, as WLAN coverage cannot be extended beyond the backhaul deployment. Second, centralized structures work inefficiently with new applications, such as wireless gaming, requiring peer-to-peer connectivity. Third, a fixed topology inhibits stations from choosing a better path for communication. WLANs can benefit significantly if they evolve to address these emerging needs.

Wireless mesh networks (WMNs) hold the promise of a solution. However, existing WMNs rely on the IP layer to enable multihop communication and do not provide an inherently wireless solution. Since wireless links are less reliable than wired links, a multihop routing protocol operating in a wireless environment must account for the nature of the wireless links. As 802.11 does not specify the interfaces that the IP layer needs to derive link metrics from the medium access control (MAC) layer, the ad hoc routing protocols developed in the Internet Engineering Task Force's (IETF's) Mobile Ad Hoc Networks (MANET) group are forced to rely on indirect measurements [1] to observe the radio environment. However, the acquired link metrics are of limited accuracy [2], whereas the MAC layer has adequate knowledge of its radio neighborhood to make its measurements less outdated and more precise. Furthermore, for transparent support of important protocols like Address Resolution Protocol, Dynamic Host Configuration Protocol, Spanning Tree, and many more, a WMN must appear like traditional LAN segments that form single broadcast domains. In encapsulating layer 2 traffic, IP-based WMNs emulate LAN behavior. However, this appears to be more like an engineering patch and does not provide a long-term solution aimed at sustainable scaling of WLANs to new applications. Since MAC-based multihop solutions inherently support layer 2 traffic, they operate transparently to any higher-layer protocol.

To realize the benefits a MAC-based WMN promises, an integrated mesh networking solution is under development in IEEE 802.11 Task Group S. The particular amendment of the 802.11 standard dealing with mesh support, 802.11s [3], describes a WMN concept that intro-

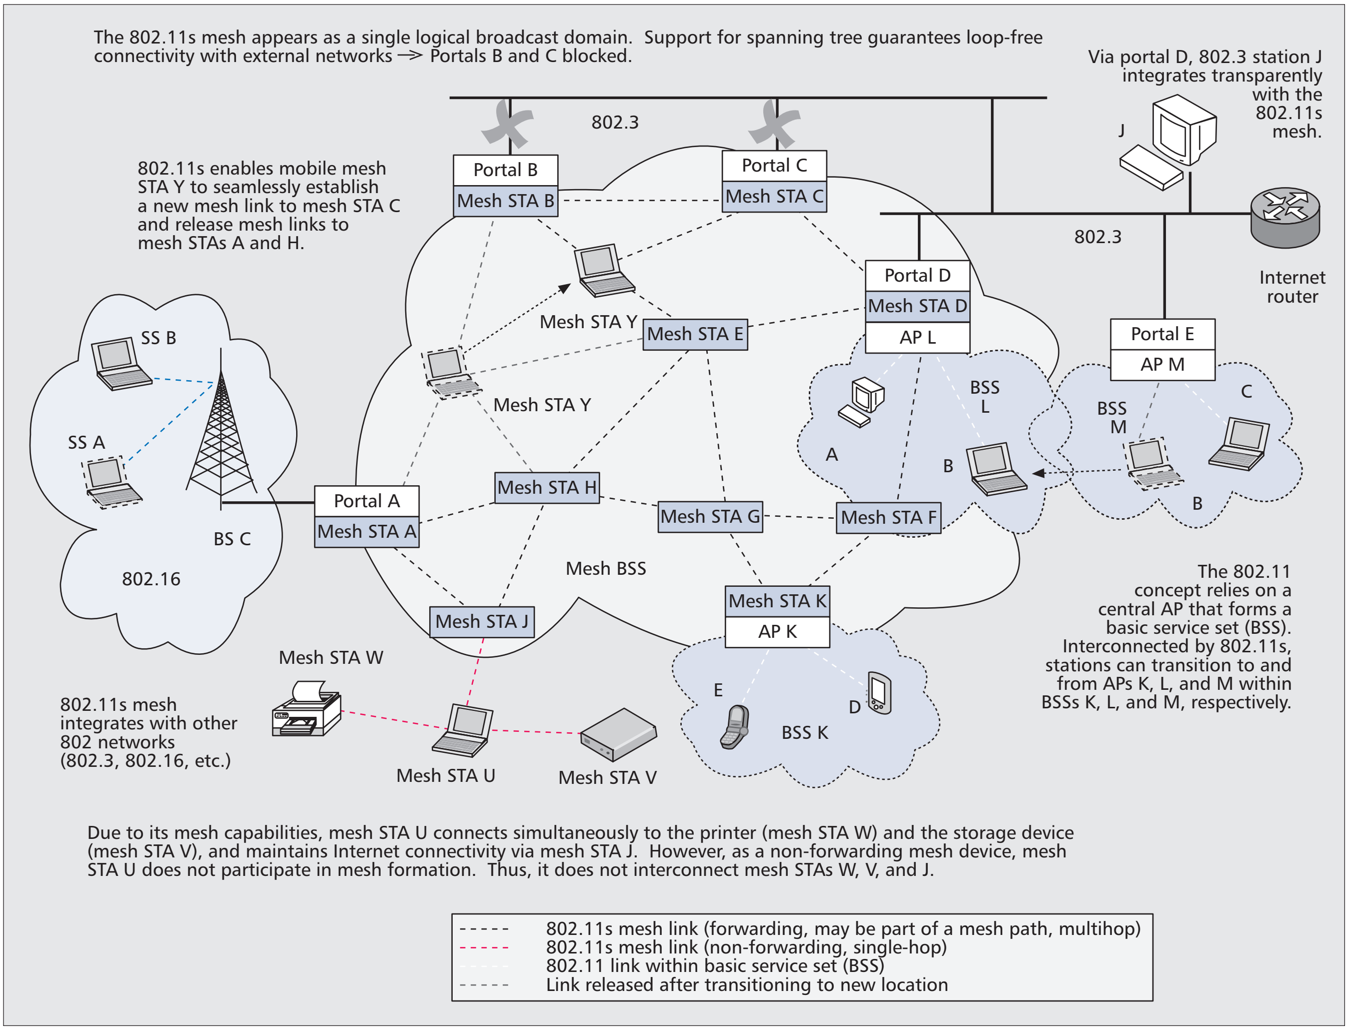

The 802.11s mesh appears as a single logical broadcast domain. Support for spanning tree guarantees loop-free connectivity with external networks → Portals B and C blocked.

Via portal D, 802.3 station J integrates transparently with the 802.11s mesh.

802.11s enables mobile mesh STA Y to seamlessly establish a new mesh link to mesh STA C and release mesh links to mesh STAs A and H.

The 802.11 concept relies on a central AP that forms a basic service set (BSS). Interconnected by 802.11s, stations can transition to and from APs K, L, and M within BSSs K, L, and M, respectively.

802.11s mesh integrates with other 802 networks (802.3, 802.16, etc.)

Due to its mesh capabilities, mesh STA U connects simultaneously to the printer (mesh STA W) and the storage device (mesh STA V), and maintains Internet connectivity via mesh STA J. However, as a non-forwarding mesh device, mesh STA U does not participate in mesh formation. Thus, it does not interconnect mesh STAs W, V, and J.

Figure 1. *802.11s enables seamless connectivity among dissimilar 802 networks.*

duces routing capabilities at the MAC layer. *Path selection* is used to refer to MAC-address-based routing and to differentiate it from conventional IP routing.

To understand the general WLAN concept, we begin by explaining the 802.11 standard [4] briefly. Next, we provide an outline of the basic building blocks of 802.11s, extending what is explained in [5], wherein we discuss interworking, MAC, security, and path selection. Subsequently, we discuss certification activities for mesh solutions in the Wi-Fi alliance [6]. The last section is dedicated to implementations such as One Laptop Per Child (OLPC) [7] and the open80211s [8] project, which reveal the performance of the 802.11s draft and provide us with a preview of what can be expected with 802.11s-certified products.

The 802.11 Network Design

A station, or STA, is an 802.11-standard-compliant MAC and physical layer (PHY) implementation [4] and constitutes the basic entity in an 802.11 network. The most elementary 802.11 network, called a basic service set (BSS), can be formed using two stations. If a station provides the Integration service to the other stations, this station is referred to as an access point (AP). If an AP is present in a BSS, it is referred to as an infrastructure BSS. To join an infrastructure BSS, a station associates with the AP. Figure 1 provides an example where AP M is part of the infrastructure. AP M provides stations B and C with access to the distribution system (DS). The DS provides the services that are necessary to communicate with devices outside the station's own BSS. Furthermore, the DS allows APs to unite multiple BSSs to form an extended service set (ESS). Within an ESS, stations can roam from one BSS to another [9]. Today Ethernet (802.3) usually provides the distribution system medium (DSM) on which the DS relies. Consequently, in practice, APs collocate with the so-called portals that provide the integration of WLANs with non-802.11 networks.

The IEEE 802.11 standard itself does not provide any details about the DSM. In principle, the DSM can be wireless too. The 802.11 frame format (without the extensions highlighted in Fig. 2) provides four fields necessary for addressing over multiple intermediate devices. The source address indicates the station that generated the frame (initial hop), and the destination address indicates the intended receiver (final hop). Both addresses remain unchanged in a concatenated set of multiple wireless hops. The

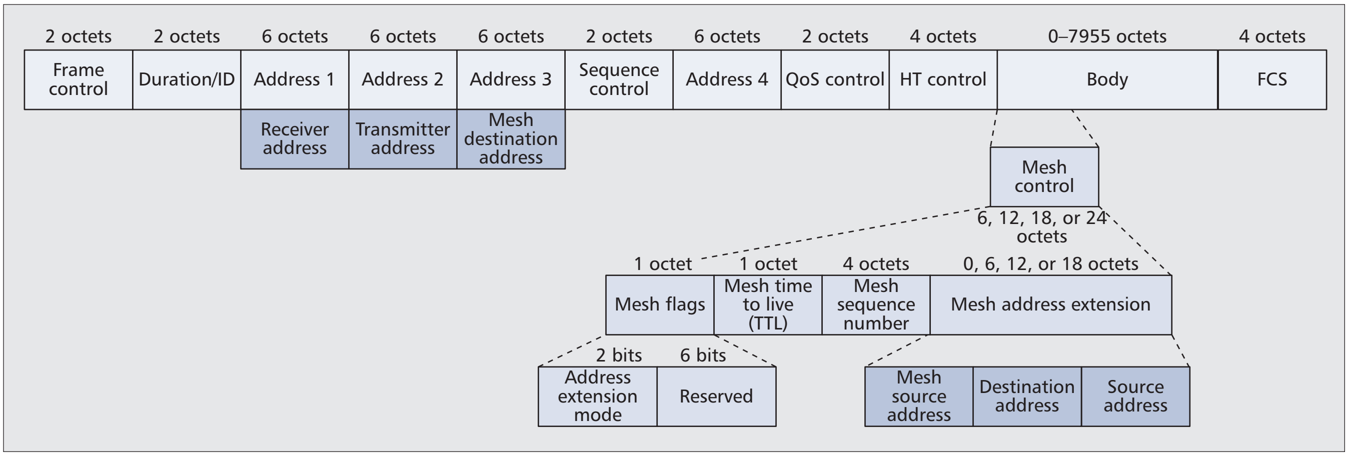

Figure 2. *The 802.11s mesh control field is part of the frame body and provides up to two more address fields.*

transmitting and receiving station addresses, which denote the stations that actually forwarded the frame, change with every hop. The 802.11 frame format provides two additional bits denoted "To DS" and "From DS." The bit combinations 10 and 01 indicate the traffic entering or leaving the DS from a BSS, respectively. For the traffic that is relayed within the DS from one AP to another, the bit combination 11 is used.[1] Interestingly, vendors often use the vernacular term wireless DS (WDS) for this configuration.

The 802.11s Concept

Since the current standard [4] explicitly states that it does not define the procedures necessary for WDS implementation, many 802.11 multihop implementations do not interoperate. 802.11s not only helps interconnect BSSs wirelessly and thereby fills the WDS gap; it also enables a new type of BSS, the so-called mesh BSS (MBSS). In the following we describe the major sections of 802.11s. Its amendments to interworking, MAC, security, and path selection make the MBSS a self-contained network that enables applications beyond what a traditional single-hop WLAN supports.

Interworking

As a family of standards, interoperability between the different networking concepts is a requirement for 802. For seamless integration, the 802.11s network appears as a single Ethernet segment to the outside (Fig. 1). The WMN implements a single broadcast domain and thus integrates seamlessly with other 802 networks. In particular, 802.11s supports transparent delivery of uni-, multi-, and broadcast frames to destinations in- and outside of the MBSS (referred to as mesh in the following). Devices that form the mesh are called mesh stations (mesh STAs). Mesh stations forward frames wirelessly but do not communicate with non-mesh stations. However, a mesh station may be collocated with other 802.11 entities.

Frame Structure

Currently, 802.11 categorizes frames as data, control, or management. Data frames carry higher-layer data. Control frames are used for acknowledgments and reservations. Devices use management frames to set up, organize, and maintain a WLAN and the local link. To provide for multihop, 802.11s extends data and management frames by an additional mesh control field, as shown in Fig. 2. The mesh control field consists of a mesh time to live (TTL) field, a mesh sequence number, a mesh flags field, and possibly a mesh address extension field. The TTL and sequence number fields are used to prevent the frames from looping forever. When mesh stations communicate over a single hop, their frames do not carry the mesh control field.

The mesh flags field indicates the presence of additional MAC addresses in the mesh control field. The address extension allows for a total of six address fields in a mesh frame. This is useful when the source and destination of the frame are not part of the mesh, but are proxied by mesh stations. Figure 1 presents an example where mesh station D proxies non-mesh stations A, B, and J. Informing other mesh stations of its proxied devices, mesh station D diverts to itself all frames destined for A, B, or J. Together with the six-address scheme, the proxied entities can be identified as the final destination beyond the intermediate destination D. In addition, the extension to six addresses allows for proactive routing, explained later. Proactive routing divides a path into two distinct routes to simplify path selection. In Fig. 3 only mesh station C maintains paths to all mesh stations. In this case non-mesh station D's frames enter the mesh at mesh station K, traverse to mesh station C (the first route), and from there to mesh station J (the second route). An observant reader will note from Fig. 2 that the address extension field allows for the addition of three addresses, rather than just two. The rationale for this is that standard management frames have three addresses only. Hence, in the case of multihop mesh management frames, address 4 is included in the mesh control field rather than in the standard frame header.

Mesh Formation and Management

Just as an AP's beacon frame helps the stations to detect a BSS and learn about its settings, the mesh station's beacon carries information about the mesh and helps other mesh stations detect and join the mesh. Mesh stations detect each

[1] *The bit combination 00 is to be used within an independent BSS that does not have an AP.*

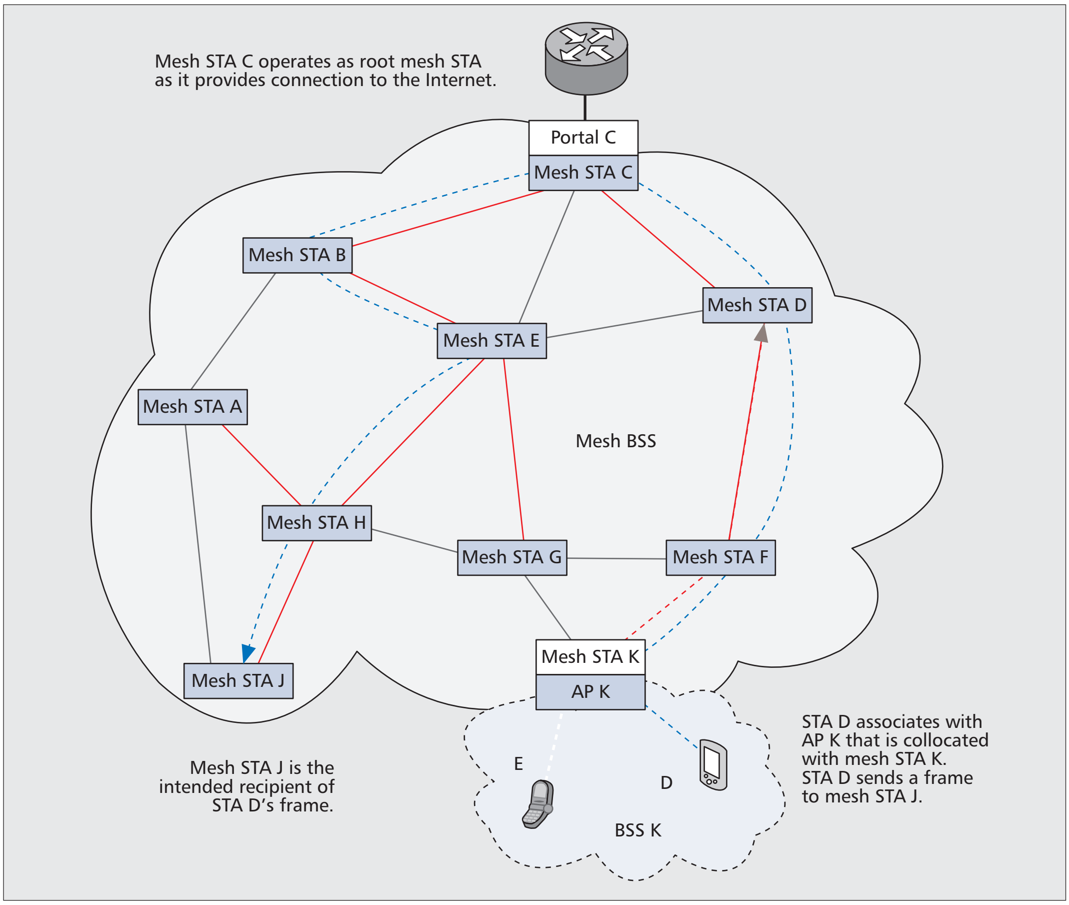
Mesh STA C operates as root mesh STA
as it provides connection to the Internet.
Portal C
Mesh STA C
Mesh STA B
Mesh STA D
Mesh STA E
Mesh STA A
Mesh BSS
Mesh STA H
Mesh STA G
Mesh STA F
Mesh STA K
AP K
Mesh STA J
E
D
BSS K
STA D associates with AP K that is collocated with mesh STA K. STA D sends a frame to mesh STA J.
Mesh STA J is the intended recipient of STA D's frame.

Figure 3. *The six address scheme provides support for proxied stations and tree-based path selection.*

If a mesh station holds a larger channel precedence value, it broadcasts its value and may indicate a different frequency channel. Following the highest announced precedence value, mesh stations finally coalesce on the new channel.

other based on passive scanning (observation of beacon frames) or active scanning (probe frame transmission). The mesh-specific beacon and probe frames contain a Mesh ID (the name of a mesh), a configuration element that advertises the mesh services, and parameters supported by the transmitting mesh station. This functionality enables mesh stations to search for suitable peers (e.g., other mesh stations that use the same path selection protocol and metric). Once such a candidate peer has been identified, a mesh station uses the Mesh Peer Link Management protocol to establish a peer link with another mesh station. Even when the physical link breaks, mesh stations may keep the peer link status to allow for quick reconnection. In Fig. 1 mesh station Y may re-establish connection with mesh station A or H as soon as it moves in range again.

Mesh stations use a single transceiver only. Accordingly, a mesh operates in a single frequency channel only. With multitransceiver devices, however, different frequency channel meshes can be unified into a single LAN. Figure 4 provides an example where five meshes operate in four different frequency channels. Mesh stations C, D, and E collocate within a device that has three independent transceivers. Incorporating an 802 bridge in the device, the collocated mesh stations interconnect and help to forward frames between their meshes. Consequently, a single WMN can be constituted.

Regulatory bodies have different requirements on the frequency bands used by 802.11 and mesh stations are required to comply with these regulatory requirements. In Europe, for example, devices must switch to a different channel (dynamic frequency selection) upon detection of a radar station in the 5 GHz band. To prevent a mesh from splitting, a channel selection protocol allows selection of the new frequency channel. In the absence of a central coordinator, a distributed algorithm is developed, based on a 31-bit random channel precedence value, for arbitration.

If mesh station O in Fig. 4 detects a radar station, it is required to leave its current frequency channel and indicates the new frequency channel to mesh station N. N forwards the message, and thus, mesh stations D and T learn about the channel switch too. After a predetermined period of time, the mesh channel switch time, the mesh stations switch to the new channel. If a mesh station holds a larger channel precedence value, it broadcasts its value and may indicate a different frequency channel. Following the highest announced precedence value, mesh stations finally coalesce on the new channel.

Synchronization and Power Management

All beacon frames provide a time reference that is used for synchronization and power saving. Power-saving mesh stations are either in light- or deep-sleep mode. Being in light-sleep mode, a mesh station switches to full power whenever a neighbor or the mesh station itself is expected to transmit a beacon frame. In deep-sleep mode a

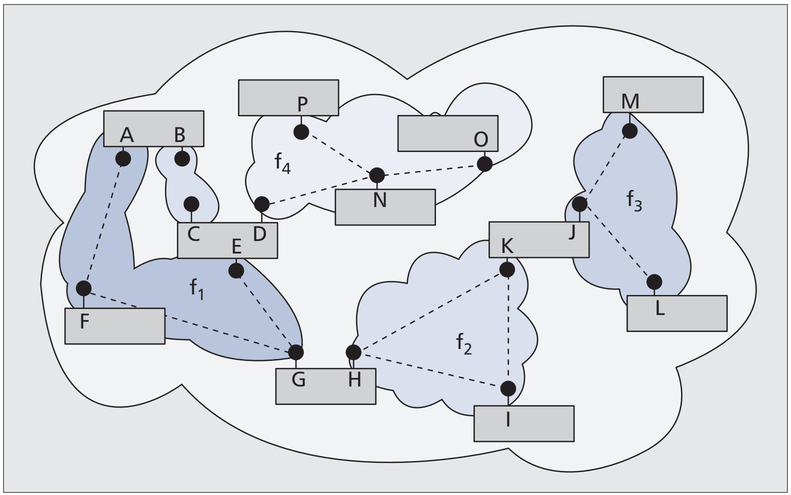

Figure 4. *Layer 2 bridges in multi-transceiver devices may unify different frequency channel meshes into a single LAN.*

mesh station solely wakes up for its own beacon frame transmissions. The mesh station can be informed of buffered traffic during the awake period that follows the beacon. Synchronization enables a new kind of distributed reservation protocol, introduced in the next section.

Medium Access Control in 802.11s

For medium access, mesh stations implement the mesh coordination function (MCF). MCF consists of a mandatory and an optional scheme. For the mandatory part, MCF relies on the contention-based protocol known as Enhanced Distributed Channel Access (EDCA), which itself is an improved variant of the basic 802.11 distributed coordination function (DCF). Using DCF, a station transmits a single frame of arbitrary length. With EDCA, a station may transmit multiple frames whose total transmission duration may not exceed the so-called transmission opportunity (TXOP) limit. The intended receiver acknowledges any successful frame reception. Additionally, EDCA differentiates four traffic categories with different priorities in medium access and thereby allows for limited support of quality of service (QoS).

To enhance QoS, MCF describes an optional medium access protocol called Mesh Coordinated Channel Access (MCCA). It is a distributed reservation protocol that allows mesh stations to avoid frame collisions. With MCCA, mesh stations reserve TXOPs in the future called MCCA opportunities (MCCAOPs). An MCCAOP has a precise start time and duration measured in slots of 32 μs. To negotiate an MCCAOP, a mesh station sends an MCCA setup request message to the intended receiver. Once established, the mesh stations advertise the MCCAOP via the beacon frames. Since mesh stations outside the beacon reception range could conflict with the existing MCCAOPs, mesh stations also include their neighbors' MCCAOP reservations in the beacon frame. At the beginning of an MCCA reservation, mesh stations other than the MCCAOP owner refrain from channel access. The owner of the MCCAOP uses standard EDCA to access the medium, and does not have priority over stations that do not support MCCA. Although this compromises efficiency, simulations reveal that high medium utilization can still be achieved with MCCA in the presence of non-MCCA devices [10]. After an MCCA transmission ends, mesh stations use EDCA for medium contention again.

Congestion Control

Access in 802.11 relies on carrier sensing. At a mesh's edge, mesh stations have fewer neighbors. and therefore observe an idle wireless medium more often than mesh stations located in the core of the mesh. Consequently, edge mesh stations have a higher probability to transmit. When core mesh stations congest, they cannot carry the aggregated traffic and drop frames. This is costly as the mesh frame has already traversed several hops to reach the congested mesh station. The optional 802.11s congestion control concept uses a management frame to indicate the expected duration of congestion and to request a neighbor mesh station to slow down. Since it is each mesh station's choice to issue a congestion control frame, the notification may finally ripple back to the traffic source.

Security in 802.11s

With 802.11s, mesh stations perform the dictionary attack-proof Simultaneous Authentication of Equals (SAE) [11] algorithm. Besides mutual authentication, SAE provides two mesh stations with a pairwise master key (PMK) that they use to encrypt their frame. As its name indicates, SAE does not rely on a keying hierarchy like traditional 802.11 encryption [4]. Instead, it implements a distributed approach that both mesh stations may initiate simultaneously. Because of the pairwise encryption, each link is independently secured. As a consequence, 802.11s does not provide end-to-end encryption. Since broadcast traffic must reach all authenticated peers, a mesh station is required to update its broadcast traffic key with every new peering it establishes.

Path Selection in 802.11s

Within a mesh, all mesh stations use the same path metric and path selection protocol. For both, 802.11s defines a mandatory default scheme. Because of its extensible framework, they can be replaced by other solutions.

The default metric, termed airtime metric, indicates a link's overall cost by taking into account data rate, overhead, and frame error rate of a test frame of size 1 kbyte. The default path selection protocol, Hybrid Wireless Mesh Protocol (HWMP), combines the concurrent operation of a proactive tree-oriented approach with an on-demand distributed path selection protocol (derived from the Ad Hoc On Demand Distance Vector [AODV] protocol [1]). The proactive mode requires a mesh station to be configured as a root mesh station. In many scenarios this will be a mesh station that collocates with a portal (Fig. 3). As such, the root mesh station constantly propagates routing messages that either establish and maintain paths to all mesh stations in the mesh, or simply enable mesh stations to initiate a path to it (red lines in Fig. 3). In the example of Fig. 3, mesh station K uses the root mesh station C to establish an initial path (dotted line) to mesh station J. Once established, mesh stations may use the AODV part of HWMP to avoid the indirection via the root mesh station. In the present example, K could discover a shorter path (links marked in grey) via G and H to forward station D's frames to the destination mesh station J. Mesh stations also rely on AODV when a root mesh station is unavailable. When no path setup messages are propagated proactively, however, the initial path setup is delayed.

To allow for even simpler configuration, vendors may opt not to implement HWMP at all. As an example, a battery-limited handheld device may refrain from frame forwarding to minimize power consumption. Accordingly, it does not propagate path information and behaves like an end station. However, the device is still able to request the frame forwarding service from neighboring mesh stations.

Wi-Fi Alliance's Mesh Marketing Task Group

In October 2006 the Wi-Fi Alliance (WFA) established the Mesh Marketing Task Group, chartered to work on a Marketing Requirement Document, and the specification of a certification and test plan. To meet customers' expectations, WFA's mesh activities aim at compliance with the present certification programs. Accordingly, simple and secure setup of a WFA-certified mesh needs to comply with the existing WFA programs. The current Wi-Fi protected setup already enables security via a push button, secret PIN, or near-field communication, for example.

To provide compatibility with existing Wi-Fi devices, WFA's marketing program requires each mesh station to incorporate either the AP or station functionality too. While Wi-Fi mesh APs must support frame-forwarding and thereby help to increase the radio coverage, non-AP mesh stations may choose to become an end station. Whereas a non-mesh station connects to a single AP only, a mesh station may connect with multiple other mesh stations even if it does not forward traffic for others. Consequently, it provides users the advantage of access to services not reachable via the AP (mesh station U in Fig. 1).

Currently Deployed Implementations

The OLPC project and open80211s [7] are the world's first implementations of 802.11s. In the next two sections we briefly introduce the project goals and experiences gained from real-world setups.

The OLPC Project

Developed by the OLPC Foundation, the XO laptop aims to serve as a learning tool for children living in developing countries where a communication infrastructure is unlikely to exist. With WLAN embedded in the XO, the decision to implement 802.11s was self-evident. Based on an early draft of 802.11s, the XO omits certain functions of 802.11s such as encryption or proactive routing. One of the challenges faced by OLPC was to ensure at all times a minimal node density, which is critical for the proper operation of a mesh network. Two design choices were made to address this issue [7].

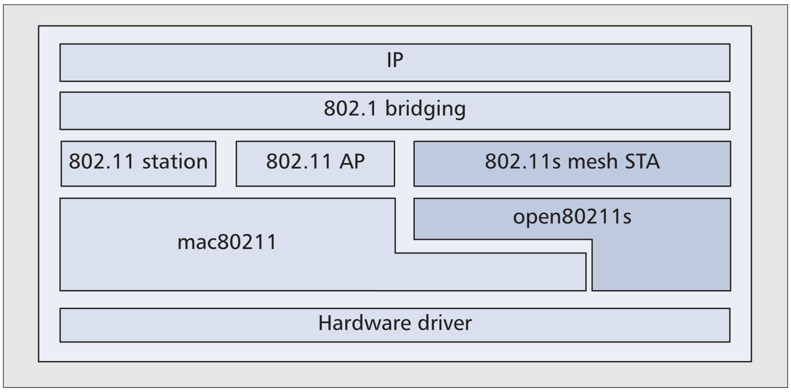

Figure 5. *The open80211s stack integrates into the Linux kernel.*

First, the OLPC mesh does not implement any access control mechanism. Each node can receive and forward traffic from any other mesh-capable node, thus avoiding a possible fragmentation of the network caused by incompatible access credentials. As there is no authentication at the mesh layer, XOs must rely on upper layers for confidentiality.

Second, the mesh protocol stack is embedded in the wireless network card. With this architecture, the entire 802.11s code can operate independently of the host CPU. Thus, the XO works as a mesh station even when in power-save mode; that is, a laptop transitioning into power-save mode will not adversely affect other students who may rely on a single student's provision of Internet connectivity.

Due to its distributed nature, OLPC assumes that a root mesh station is never available. Thus, the XO would not benefit from implementing a tree structure. Consequently, the XO solely implements HWMP's AODV part.

Several presentations at road shows and live demonstrations have shown the capabilities of the present OLPC mesh implementation. To date, OLPC has shipped over 1.2 million mesh-capable laptops around the globe.

open80211s

open80211s [7] is a vendor-neutral implementation of 802.11s for the Linux operating system (Fig. 5). Since 802.11s introduces only minimal changes to the MAC layer, the 802.11s stack can be almost fully implemented in software and made to run on legacy 802.11 cards. The goal of the project is to closely track the 802.11s draft and support the interoperability of different 802.11s implementations. The availability of the source code helps to identify and resolve design problems, and resolve ambiguities in the protocol being specified. Performance measurements are routinely taken before each release. Figure 6a shows the path discovery time for different path lengths, where we measure a 12-node open80211s testbed wherein all the nodes are in radio range

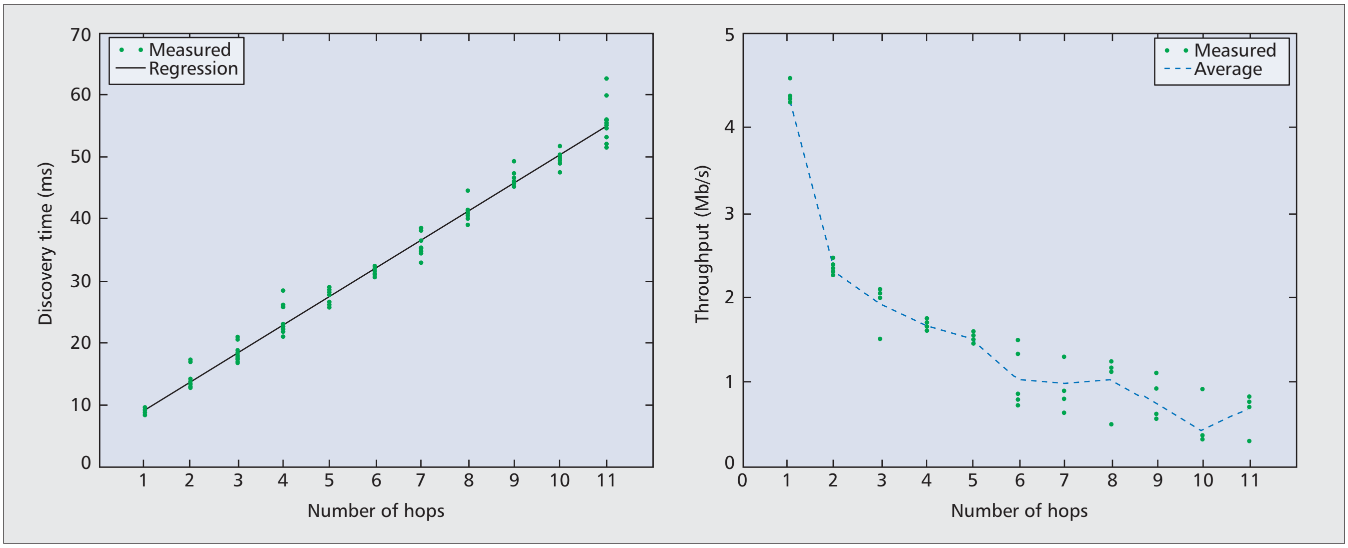

Figure 6. *open80211s performance measurements.*

of each other, and manual address filter settings enforce multihop topologies. We measure the discovery duration from the time a path setup is requested until the time the path is actually established. As expected, routes are resolved in linear time, and the variance increases with the length of the path. Figure 6b shows the average end-to-end data rate measurement results in an environment where the default 802.11 parameters are used for measurements, and where each experiment is run six times for 10 s per measurement in an uncontrolled wireless environment.

As previously mentioned, non-802.11s products rely on the unspecified WDS to enable multihop networking. However, without path selection mechanisms, suboptimal path length or undesired increase in path length occurs. Figure 6 illustrates scenarios with more than four to five hops that are usually beyond the application scenario of a single-channel WMN. Note that throughput decreases rapidly with the number of hops in the fully connected topology of the experiment. Under normal conditions, the routing protocol used in 802.11s would resolve paths that are limited to one or two hops only. The open80211s stack has been part of the Linux kernel since version 2.6.26 (July 2008).

Conclusions

The wired Internet liberated users from dial-up connections and leased lines. By interconnecting autonomous systems, the decentralized Internet enabled new services and business. With 802.11s, a similar development has started. The classical single-hop connectivity no longer addresses the needs of an ever-growing user base, and mesh technology is the natural evolution.

However, with mesh networking, 802.11 enters uncharted territories, and 802.11s does not yet define the definite solution. Further improvements are necessary to increase efficiency and enable a degree of QoS that users are willing to accept. Judging from the current status of the ongoing standardization process, it seems that the finalization of the 802.11s standard may be expected next year. Remaining problem areas are congestion control, channel selection, and medium access.

The congestion control mechanism requires a mesh station to specify the congestion duration to its neighbors. However, due to varying radio conditions, the link speed changes; a congested mesh station can hardly predict this. Even worse, legacy devices in overlapping BSSs cannot comply with congestion messages. As a result, neighboring BSSs receive an increased share of the wireless medium, and the congested mesh station is not really helped. Finally, conditions that trigger congestion control, as well as neighbor mesh station reaction, are left unspecified.

Another concern addresses the current 802.11s method for distributed frequency channel selection. Based on random values that are needed for arbitration, the propagation of a common frequency channel has limited reliability. Unfortunately, the selection scheme cannot guarantee that all mesh stations are informed. Since a mesh station switches at the latest after 255 ms, large networks are subject to partitioning when different mesh stations increase the precedence value and the message cannot propagate back to the originator in time.

Finally, there remains the challenge of medium access. Currently, WLAN deployments rely on a wired backbone where APs need not be in radio range and hence do not share a frequency channel. However, to form a WMN, this changes. With the wireless medium being shared among neighbors, the environment becomes much more interference-prone. Furthermore, mesh stations have no priority over other 802.11 devices, and a mesh suffers severely from any overlapping BSS. Even if a mesh AP incorporates multiple transceivers to separate the BSS and mesh network into different frequency bands (2.4 and 5 GHz), the WMN carries the aggregation of locally generated and forwarded traffic. Accordingly, the WMN is threatened with saturation. Experiments with real 802.11s deployments substantiate these limitations of the EDCA-based medium access mechanisms when used in a WMN environment. Schemes tailored to mesh-specific needs, such as the MCCA scheme, possibly combined with path selection and flow control, are quite likely to bring benefits.

Aided by the experiences gained in the implementation projects, IEEE 802.11s is in the process of tackling these challenges. Once overcome, mesh technology will be an inherent part of any future wireless standard. Today a novelty, users are likely to take the ability to communicate without a wired infrastructure for granted in the future: convenient deployment and spontaneous connectivity — anytime, anywhere. To provide for this, 802.11s will not remain the only mesh solution. With 802.11s, the fascinating WMN adventure has just begun.

Aided by the experiences gained in the implementation projects, IEEE 802.11s is in the process of tackling these challenges. Once overcome, mesh technology will be an inherent part of any future wireless standard.

References

[1] C. Perkins, E. Belding-Royer, and S. Das, "Ad Hoc On-Demand Distance Vector (AODV) Routing," IETF RFC 3561, Jul. 2003.

[2] I. F. Akyildiz, X. Wang, and W. Wang, "Wireless Mesh Networks: A Survey," *Comp. Networks and ISDN Sys.*, vol. 47, no. 4, Mar. 2005.

[3] "Draft Standard for Information Technology — Telecommunications and Information Exchange Between Systems — LAN/MAN Specific Requirements — Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) specifications: Amendment 10: Mesh Networking," IEEE unapproved draft, IEEE P802.11s/D4.0, Dec. 2009.

[4] IEEE P802.11-2007, "Information Technology — Telecommunications and Information Exchange Between Systems — Local and Metropolitan Area Networks — Specific Requirements — Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications," JunE 2007.

[5] S. M. Faccin *et al.*, "Mesh WLAN Networks: Concept and System Design," *IEEE Wireless Commun.*, vol. 13, no. 2, Apr. 2006.

[6] Wi-Fi Alliance, available: http://wi-fi.org/

[7] J. Cardona, "Wireless Meshing with the One Laptop Per Child," http://www.cozybit.com/whitepapers/olpc-mesh.pdf

[8] open80211s, http://open80211s.org/

[9] B. Walke, S. Mangold, and L. Berlemann, *IEEE 802 Wireless Systems: Protocols, Multi-Hop Mesh/Relaying, Performance and Spectrum Coexistence*, Wiley, Nov. 2006.

[10] Y. Chen, and S. Emeott, "MDA Simulation Study: Robustness to Non-MDA Interferers," IEEE 802 Plenary Meeting, Submission 11-07-0356-00-000s, Orlando, FL, Mar. 2007.

[11] D. Harkins, "Simultaneous Authentication of Equals: A Secure, Password-Based Key Exchange for Mesh Networks," *SENSORCOMM '08 — 2nd InGut'l. Conf. Sensor Technologies and Apps.*, Aug. 2008.

Biographies

Guido R. Hiertz (hiertz@ieee.org) received his Dipl.-Ing. degree in electrical engineering from RWTH Aachen University, Germany. Working toward his Ph.D. in the Department of Communication Networks, he contributed to various research projects and authored several papers at IEEE conferences. He is a voting member of IEEE 802.11 and Vice Chair of IEEE 802.11s. He is a charger member of the industry forum Wi-Mesh Alliance that created the initial draft of IEEE 802.11s jointly with the industry group SEE-Mesh. Since 2009 he is the head of research and development in the Rental Series Department at Riedel Communications, Wuppertal, Germany.

Dee Denteneer received an M.Sc. in statistics from the University of Utrecht and a Ph.D. in applied probability (queuing analysis) from Eindhoven University ot Technology, the Netherlands. From 1984 to 1988 he worked at the Dutch Central Statistical Office, where he designed the Blaise language, a language for questionnaire description. Since 1988 he has been employed at Philips Research in Eindhoven, since 2000 as a principal research scientist. At Philips he has worked on the application of mathematics in various industrial research projects such as MPEG encoding, speech recognition, secure biometrics, and data transmission systems. He is Chair of IEEE 802.11s. His current research interest is in the performance analysis and standardization of wireless mesh networks.

Sebastian Max studied computer science at RWTH Aachen University, and received his diploma degree with distinction in 2005. Since then he has been with the Chair of Communication Networks (ComNets) at RWTH Aachen University, where he is working toward his Ph.D. He holds the Research College (Graduiertenkolleg) Software for Mobile Communication Systems scholarship of the German Research Foundation (DFG). His main research field is wireless mesh networks for city-wide Internet access.

Rakesh Taori is a principal engineer in the Digital Media and Communications R&D Centre at Samsung Electronics, Suwon, South Korea. He is currently involved in research, development, and standardization of the 4G air interface technologies pertaining to the MAC layer. He was an active contributor to the standardization of multihop relay in the IEEE 802.16 systems and is now contributing to the development of the IEEE 802.16m amendment: the advanced OFDMA air interface for the IEEE 802.16 system. Prior to joining the Samsung DMC R&D Centre, he held research positions at Samsung Research (2004–2008, South Korea), Ericsson Research (2000–2004, Sweden and the Netherlands), and Philips Research Laboratories (1992–2000, the Netherlands). Over the past 17 years he has performed research and standardization work in the area of media coding and wireless systems. In the area of media coding his primary focus was low-bit-rate parametric coding of speech and audio signals. In the area of wireless systems he has contributed to research and standardization in wireless PANs (Bluetooth and UWB) and wireless LANs (802.11s, WLAN mesh), and is currently active in the area of wireless MANs (802.16m, advanced air interface). He has contributed to several standardization organizations — MPEG, ITU-T, ETSI, Bluetooth SIG, IEEE, and the WiMAX Forum — and has served these organizations in various roles. From August 2004 to November 2005 he served as Chair of the Technical Steering Committee of the WiMedia Alliance. He obtained his B.Eng. degree in control and computer engineering and M.Phil degree in digital signal processing and communications from the University of Westminster, London, United Kingdom.

Javier Cardona is the co-founder and CEO of cozybit Inc., an engineering consulting firm in the field of wireless communications. He was one of the implementers of the OLPC mesh stack and open802.11s. He holds an M.S. in telecommunications from Universitat Politecnica de Catalunya, Spain, and an M.Eng. in embedded systems design from Alari.

Lars Berlemann (lars.berlemann@t-mobile.net) is program manager at T-Mobile International, leading the market introduction program for NGMN. He holds a Ph.D. and Diploma degree in electrical engineering from RWTH Aachen University as well as a Diploma degree in business and economics from the same university. He is author of the textbooks *Cognitive Radio for Dynamic Spectrum Access* (Wiley, 2009) and *IEEE 802 Wireless Systems: Protocols, Multi-Hop Mesh Relaying, Performance and Spectrum Coexistence* (Wiley, 2006). He has published a multitude of reviewed publications including several journal articles, and was scientific organizer of European Wireless Conference 2005 and IEEE PIMRC 2005. He has been a guest lecturer on mobile radio networks at the Technical University of Dortmund since 2007.

Bernhard Walke is directing the ComNets Research Group at RWTH Aachen University, Germany, focusing on 4G air interface design and performance evaluation, besides developing system-level simulation tools like openWNS. He contributed, together with his Ph.D. students, revolutionary concepts that are being used in standardized mobile radio networks, such as the packet data traffic channel of GPRS operating on a GSM traffice channel; the fast radio link establishment for GPRS (later named TBF), a concept also used in UMTS; the MAC frame applied in WiMAX and 3GPP LTE systems for radio resource allocation; and the concept of fixed decode-and-forward relays used in broadband cellular radio networks like WiMAX and 3GPP LTE/LTE-A. Besides that, his group has contributed a number of improvements now implemented in IEEE 802 standards. His work is published in six textbooks and about 260 peer reviewed conference and journal papers. Prior to joining academia, he worked for 18 years in industry at EADS AG. He holds a doctorate in information engineering from the University of Stuttgart, Germany.