

Proposed Routing for IEEE 802.11s WLAN Mesh Networks

Michael Bahr

Siemens Corporate Technology, Information & Communications

Otto-Hahn-Ring 6

81730 München, Germany

bahr@siemens.com

ABSTRACT

This paper describes the proposed routing for IEEE 802.11s WLAN mesh networks based on the current draft standard D0.01 from March 2006. IEEE 802.11s defines a new mesh data frame format and an extensibility framework for routing. The default routing protocol HWMP is described. HWMP is based on AODV and has a configurable extension for proactive routing towards so-called mesh portals. It uses MAC addresses (layer 2 routing) and uses a radio-aware routing metric for the calculation of paths. Furthermore, the optional routing protocol RA-OLSR is described.

Note, that the standardization of WLAN Mesh Networking in IEEE 802.11s is work in progress during the time of writing. While the general concepts of the proposed routing protocols seem to be quite fixed, the details are likely to change.

Categories and Subject Descriptors

C.2.2 [Computer-Communication Networks]: Network Protocols – *Routing protocols.*

General Terms

Algorithms, Standardization

Keywords

Wireless Mesh Networks, Routing, IEEE 802.11

1. INTRODUCTION

Wireless mesh networks have received increased attention over the last years. The number of installations of wireless mesh networks (WMNs) is increasing continuously. Several successful start-up companies exist, so-called “mesh companies.” They have been around for several years, but now they are selling mesh products, provide wireless mesh solutions to customers, and their names are well established. The increasing number of publications

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

WICON'06, The 2nd Annual International Wireless Internet Conference, August 2-5, 2006, Boston, MA, United States

© 2006 ACM 1-59593-514-2/06/08...\$5.00

and of press articles related to wireless mesh networks raised the awareness of and increased the publicity for this kind of wireless networks.

Another indicator for the increased interest in wireless mesh networks are the several new standardization groups for WMNs and the large interest in them. IEEE 802.11s standardizes WLAN mesh networks. IEEE 802.15.5 works on mesh networking for wireless personal area networks. IEEE 802.16j defines wireless multi-hop relaying

Wireless mesh networks promise greater flexibility, increased reliability, and improved performance over conventional wireless LANs. The main characteristic of wireless mesh networking is the communication between nodes over multiple wireless hops on a meshed network graph. Efficient routing protocols provide paths through the wireless mesh and react to dynamic changes in the topology, so that mesh nodes can communicate with each other even if they are not in direct wireless range. Intermediate nodes on the path will forward the packets to the destination

Mobile ad hoc networks (MANETs) are based on the same principles – wireless multi-hop communication and efficient routing protocols for wireless meshed network graphs. In fact, the routing protocols developed for MANETs are often applied to wireless mesh networks.

Wireless mesh networks and mobile ad hoc networks use the same key concepts, but they emphasize different aspects. MANETs evolved from an academic environment and focus on end user devices, mobility, and ad hoc capabilities. In contrast to this, WMNs come from a business background and focus on mainly static devices, often infrastructure devices, reliability, network capacity, and, of course, practical deployment. Nevertheless, there is no strict border between MANETs and WMNs. Both terms can be found together in articles or publications, indicating their close relation.

The most prominent usage scenario of wireless mesh networks is currently public wireless access. The wireless mesh network provides a flexible backhaul for WLAN access points, which are distributed throughout cities or university and company campuses.

A survey on wireless mesh networks can be found in [3]. An overview on routing in WMNs is given in [5]. The proposed routing in the upcoming IEEE 802.11s standard on WLAN mesh networking is described in this paper. The paper is based on the current draft standard D0.01 from March 2006 [1].

Note: The standardization of WLAN mesh networks in IEEE 802.11s is work in progress during the time of writing. The task group “s” is actively working on improving the draft standard. Many comments will be expected during the first letter ballot, which will change the draft standard. Nevertheless, the general concepts of the proposed routing protocols of IEEE 802.11s seem to be quite fixed. However, changes in the details are likely. Therefore, the information of this paper should only be used with appropriate care.

The remainder of the paper is structured as follows. Section 2 gives a brief overview of IEEE 802.11s and introduces special terminology. Section 3 describes the new data frame format. The Hybrid Wireless Mesh Protocol, the default routing protocol, is explained in detail in section 4. The optional Radio-Aware Optimized Link State Routing protocol is described in section 5. Section 6 defines the proposed radio-aware link metric. The extensibility framework is explained in section 7. An outlook concludes the paper.

2. OVERVIEW OF IEEE 802.11S

The study group for ESS mesh networking of the IEEE 802.11 working group became task group “s” (TGs) in July 2004. Its goal is the development of a flexible and extensible standard for wireless mesh networks based on IEEE 802.11 [2]. One of the key functionalities of IEEE 802.11s is the wireless multi-hop routing, which sets up the paths for the wireless forwarding. The PAR document [8] defines the scope and certain requirements of IEEE 802.11s.

Mesh nodes are called *mesh points (MPs)* in IEEE 802.11s. A mesh point is an IEEE 802.11 station that has also mesh capabilities. Mesh capabilities means that it can participate in the mesh routing protocol and forwards data on behalf of other mesh points according to the proposed 802.11s amendment. The network cloud in Figure 1 is the mesh network and comprises all mesh points.




Figure 1. Example of an IEEE 802.11s WLAN mesh network with mesh points (MP), mesh access points (MAP), mesh portal, and non-mesh IEEE 802.11 stations (STA)

Mesh points that have additionally access point functionality are called *mesh access points (mesh APs or MAPs)*. IEEE 802.11 stations that do not have mesh capabilities can connect to mesh APs in order to send data over the mesh network (cf. Figure 1).

This also provides backward compatibility with existing conventional IEEE 802.11 stations. If the paper speaks of (*conventional*) *IEEE 802.11 stations (STAs)*, these non-mesh capable WLAN devices are meant.

A mesh point that has a connection to a wired network and can bridge data between the mesh network and the wired network is called *mesh portal (MPP)*.

Figure 2 illustrates the relation between the different (mesh) node types.




Figure 2. Relation between different IEEE 802.11 (mesh) nodes

The target size of an IEEE 802.11s WLAN mesh network is up to 32 mesh points according to [8]. However, this number should not be taken as a strict limit. It only says that a solution for large wireless mesh networks with several hundreds of mesh points is not required by IEEE 802.11s. In practice, IEEE 802.11s should be able to handle networks with up to ca. 50 mesh points.

The mesh data frames use an extension of the four address frame format which has been specified for the wireless distribution system (WDS) in the IEEE 802.11 standard [2]. This means, that IEEE 802.11s WLAN mesh networks are located in a different conceptual area than mobile ad hoc networks that use the IEEE 802.11 ad hoc mode in an independent basic service set [2].

The routing is on layer 2. The routing protocol uses MAC addresses and a radio-aware routing metric. It provides mesh unicast, multicast, and broadcast data delivery. In order to make the difference to routing on layer 3 with IP addresses more distinct, the preferred term for routing is *path selection* in IEEE 802.11s. The mesh routing architecture is extensible. This gives IEEE 802.11s mesh networks the flexibility to adapt to different usage scenarios by using routing protocols that are specialized and optimized for the anticipated scenario. IEEE 802.11s will support devices with a single radio as well as devices with multiple radios.

IEEE 802.11s will amend the MAC but changes to the PHY layer are not required. It is also compatible with higher layer protocols. Mesh security is based on IEEE 802.11i.

IEEE 802.11s mesh networks will be applicable to a wide range of usage scenarios. Four important groups of usage scenarios have been identified [7]. These usage scenarios are covered by IEEE 802.11s. The four usage scenarios are:

- *Residential* for wireless home networks,

- *Office* for wireless communication in office environments,
- *Campus/Community/Public Access* for (large scale) wireless Internet access in cities or on campuses,
- *Public Safety* for the flexible, ad hoc setup of wireless communication networks for emergency staff.

The IEEE 802.11s amendment can be split up into four major parts – routing, MAC enhancements, security, and general topics. This paper considers only the proposed routing in IEEE 802.11s.

3. FRAME FORMATS

The IEEE 802.11s amendment defines a new mesh data frame format (Figure 3). This MAC frame format is used for transmitting data within the WLAN mesh network. It is an extension of the existing data frame format with a mesh specific control field.




Figure 3. IEEE 802.11s mesh data frame format

The *frame control field* contains amongst other control information the type and subtype for the mesh data frame and the two flags *to DS* and *from DS*. The two flags are set to 1 in order to indicate that the data frame is in the wireless distribution system and therefore in the mesh network.

The four address fields contain 48-bit long MAC addresses. *Address 1* is the *receiver address* which defines the mesh point that has to receive the wireless transmission. *Address 2* is the *transmitter address* which defines the mesh point that sent this wireless data frame. *Address 3* is the *destination address* which defines the final (layer 2) destination of this data frame. *Address 4* is the source address which defines the (layer 2) source of this data frame.

The 3-byte long *mesh forwarding control field* contains two fields (cf. Figure 3). The 16-bit long *mesh end-to-end sequence number* is used to control broadcast flooding and to enable ordered delivery of mesh data frames. The mesh e2e sequence number uniquely identifies frames from a given source mesh point. The mesh end-to-end sequence number is set by the source mesh point and is kept unchanged during forwarding of the mesh data frame. The 8-bit long *time to live field (TTL)* is used to time-out mesh data frames that might have been caught in an accidental infinite forwarding loop.

Control messages of the path selection protocol are transmitted as management frames of type action. The IEEE 802.11s amendment defines a new category *mesh management* for action management frames. The value of the *action field* defines the type of the management message. The actual message is given as IEEE 802.11 information element (cf. Figure 4).




Figure 4. IEEE 802.11s mesh management action frame format

4. HYBRID WIRELESS MESH PROTOCOL (HWMP)

The *Hybrid Wireless Mesh Protocol (HWMP)* is the default routing protocol for IEEE 802.11s WLAN mesh networking. Every IEEE 802.11s compliant device will be capable of using this path selection protocol. This allows interoperability between devices of different vendors.

As a hybrid routing protocol, HWMP contains both reactive routing components as well as proactive routing components.

The foundation of HWMP is an adaptation of the reactive Ad hoc On-demand Distance Vector routing protocol (AODV) [12] called *Radio-Metric AODV (RM-AODV)* [4]. While AODV works on layer 3 with IP addresses and uses the hop count as routing metric, RM-AODV works on layer 2 with MAC addresses and uses a radio-aware routing metric for the path selection.

A mesh point, usually a mesh portal, can be configured to periodically broadcast mesh portal announcements, which sets up a tree with the mesh portal as root of the tree. One of these mesh portals that periodically broadcast mesh portal announcements will become the designated *root mesh portal* by configuration or a selection process. Depending on the configuration of this root portal, mesh points that receive a root portal announcement register with the root portal or not (*registration mode* or *non-registration mode*). The created and maintained tree allows proactive routing towards mesh portals. This proactive extension of HWMP uses the same distance vector methodology as RM-AODV and reuses routing control messages of RM-AODV.

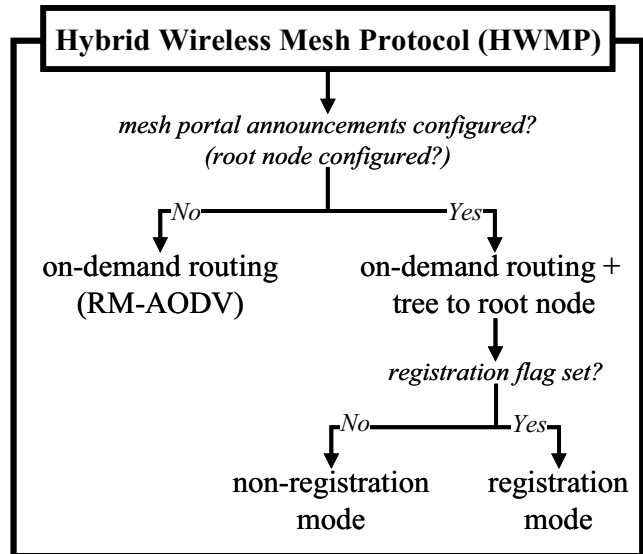


Figure 5. Configurability of HWMP

The hybrid nature and the configurability of HWMP, which is illustrated in Figure 5, provide good performance in all anticipated usage scenarios [7].

HWMP uses destination sequence numbers in order to detect outdated or stale routing information. Newly received routing information with a smaller sequence number than the sequence number of the corresponding information already known to the mesh point will be discarded, because it is outdated. This avoids the creation of routing loops and problems known from classical distance vector protocols, such as, “counting to infinity.”

Routing table entries, i.e. paths, have a lifetime associated with them. This will automatically delete unused paths when their lifetime is expired. The lifetime is reset every time data frames are transmitted over the path or by routing control messages.

4.1 Reactive Routing in HWMP

The main characteristic of reactive routing is that a path is computed only if one is needed for sending data between two mesh points. This adds an initial latency for the first packet(s) since the discovery of the links with their characteristics and the computation of the path to the requested destination start only when the first data packet has already arrived at the routing module of the source node. However, this on-demand setup of the paths uses always the most recent link state information, such as, from radio-aware link metrics and reduces the routing overhead if there is no traffic in the mesh network or the traffic pattern is not changing.

The Hybrid Wireless Mesh Protocol uses the route discovery process well-known from AODV [12] and DSR [11]. A source mesh point that needs a path to a destination mesh point broadcasts a route request message requesting a route to the destination. The route request message is processed and forwarded by all mesh points and sets up reverse paths to the originator of the route discovery. The destination mesh point or intermediate mesh points with a path to the destination will answer with a unicast route reply message. This sets up the forward path to the destination.

Furthermore, the route discovery process is adapted to the requirements of an IEEE 802.11s path selection protocol – use of layer 2 MAC addresses and use of radio aware link metrics.

The following paragraphs describe the mechanisms of the reactive routing of HWMP in more detail.

4.1.1 Generation of Route Request Messages

When a source mesh point S wants to send data to a destination mesh point D , the mesh point S first checks in its routing table whether it has a valid path to D . If not, the source mesh point S has to initiate a route discovery to D . Mesh point S is also called the *originator* of this route discovery.

S creates a *route request message (RREQ)*. The structure of the RREQ information element is shown in Figure 6.

The *RREQ ID field* together with the *source address field* uniquely identifies a route discovery, at least its route request part. This allows to detect duplicate reception of RREQ messages of the same route discovery at a node. The *source sequence number field* contains the current sequence number of the source node.




Figure 6. Structure of HWMP route request (RREQ) information element

HWMP allows to discover paths to multiple destinations simultaneously with a single RREQ message. The *destination count field* specifies the number of destination mesh points to be discovered. *Destination count* sequences of the fields *per destination flags*, *destination address*, and *destination sequence number* are contained in the RREQ. For this reason, the control flags for the RREQ have to be split up into two groups.

Those control flags that might have different values for different destinations in the RREQ are set for each destination separately in the corresponding per destination flags fields. These flags are the *destination only flag (DO)* and the *reply and forward flag (RF)*. If the DO flag is set (DO=1), only the destination D itself can create a route reply message as answer to this route request. This is the default behavior of HWMP. It ensures that the discovered path metric is current, since the route request and the route reply traverse the complete path and collect the current metric values. The RF-flag controls the forwarding of the RREQ message in case an intermediate mesh point generated a route reply message to a RREQ with DO=0.

Control flags that have the same value for all destinations in the RREQ are set in the *flags field*. This is the *unicast/ broadcast flag (UB)* which is set to broadcast (UB=1) by default. It has been introduced for the proactive extensions of HWMP.

HWMP uses an arbitrary routing metric, usually a radio-aware link metric, such as, the default airtime link metric described in section 6, instead of the hop count routing metric. The *hop count field* in the RREQ message provides information on the number of links in the path, but it is not used for the routing decision. Both *hop count* and *metric* are initialized with 0. The *time to live field (TTL)* defines the scope of the RREQ in number of hops.

The RREQ ID counter of the source mesh point is incremented before a new route request is generated. If the route request will be used for a route discovery, the sequence number of the source mesh point, the originator, is incremented by 1 before the generation of the route request.

4.1.2 Processing of Route Requests Messages

A RREQ sets up a reverse path to the originator of the route discovery/RREQ. Later on, the route reply message will travel along this reverse path. The route request also sets up or updates a path to the transmitter of the RREQ, which is done as the first step in the processing of the RREQ.

In the next step, the following fields of the received RREQ are updated. TTL is decremented by 1. Hop count is incremented by 1. And the current metric of the previous hop is added to the metric field which is the path metric from S to this mesh point.

If no path to the source mesh point S exists, a new one is created. The corresponding destination sequence number is taken from the source sequence number field. Hop count and path metric are taken from the corresponding (updated) fields of the RREQ. The next hop to S is the transmitter of the received RREQ.

If a path to the source mesh point S already exists, the mesh point checks whether it has to be updated. The existing path to S is updated if the sequence number of the RREQ is equal or greater than the sequence number of the existing routing table entry for the source mesh point S and the new path metric of the RREQ is better than the path metric in the corresponding routing table entry. The existing path to S is updated regardless of the value of the new path metric if the sequence number of the RREQ is greater than the sequence number of the corresponding routing table entry by at least a configurable threshold value. It is also updated if a newer RREQ, meaning with a greater RREQ ID, has been received.

In order to increase the stability of already established paths to the originators of RREQs, a form of path selection hysteresis is used. A mesh point switches not immediately to a path with a worse metric during the processing of a RREQ. This allows to continue to use the good path for a certain period of time in case the RREQ of the best path has been lost or until the best path metric is seen.

If the mesh point processing the RREQ is not the requested destination, it will broadcast the updated RREQ to all its neighboring mesh points, if a path to S has been created or updated and the TTL is greater than 0. Figure 7 shows the distribution of RREQs in an example scenario.




Figure 7. Distribution of RREQs

4.1.3 Generation of Route Reply Messages

Depending on the control setting in the RREQ, the requested destination and intermediate mesh points with a valid path to the requested destination can generate *route reply messages (RREPs)* in response to a received route request message.

A requested destination D will always be allowed to generate a RREP message to the source mesh point S . However, the RREP is only generated if a new path to the originator S of the RREQ has been created or an existing path to S has been updated due to the processing of the received RREQ.

Figure 8 shows the structure of the RREP information element.




Figure 8. Structure of HWMP route reply (RREP) information element

The requested destination D will put its MAC address and its sequence number into the corresponding fields of the RREP message. D will increment its sequence number by one only if the sequence number in the RREQ message is equal to that incremented value. Otherwise, the destination does not increment its sequence number before generating the RREP.

The *source address field* contains the MAC address of the originator of the route request and the corresponding sequence number, both taken from the RREQ. The *source count field* is set to 1. The capability for multiple pairs of source address–source sequence number is an extension for the registration mode of the proactive routing part.

The *lifetime field* contains the initial life time value for the forward route from S to D . The *metric field* and the *hop count field* are initialized to 0.

An intermediate mesh point I with a valid path to the requested destination D_i is only allowed to respond with a RREP message, if the corresponding destination only flag of the RREQ message is not set ($DO_i=0$). After processing the RREQ, mesh point I sticks together the two halves of the path from S to D by setting the metric field and the hop count field of the generated RREP message to the corresponding values from the routing table entry of the path to D .

While a response by intermediate mesh points reduces the initial latency caused by the on-demand route discovery, it bases its path selection partly on old, maybe already changed, path metric information of the subpath from I to D stored in the routing table. A radio-aware routing metric changes more frequently than the hop count metric. Therefore, it is preferable to collect and use the current values of the link metrics.

In order to receive the current path metric information eventually, the reply and forward flag (RF) has been introduced. The RF flag controls the forwarding of the RREQ in case the intermediate mesh point generated a RREP. If the RF flag is set ($RF=1$), the updated RREQ is forwarded (broadcast) by the intermediate mesh point. In this case, the destination only flag has to be set ($DO=1$) in order to avoid further RREPs by the succeeding intermediate mesh points on the way to the requested destination. The setting $DO=0$, $RF=0$ corresponds to the traditional behavior of AODV.

Whichever mesh point generated the RREP message, it is then unicast on the reverse path to the originator mesh point S .

The decisions and steps for the generation of RREPs have to be done for every single destination of the destination count destinations in the RREQ message with multiple requested destinations. If a RREP has been generated for destination D_i and the RREQ has not to be forwarded to the destination D_i in case of an inter-

mediate mesh point ($RF_i=0$), destination D_i is removed from the list of requested destinations in the RREQ. If there are any destinations left in this list after all destinations have been processed, the updated RREQ will be broadcast with the remaining destinations being requested. If no destination is left in the list of requested destinations, the RREQ will not be forwarded any further.

4.1.4 Processing of Route Reply Messages

A RREP sets up the forward path from the originator S to the destination D . It is unicast along the reverse path that has been set up by the RREQ.

First, the received RREP triggers the creation or update of a route to the transmitter of the RREP. The RREP is updated, that is, the hop count value is incremented by 1 and the link metric of the previous hop is added to the routing metric field. The forward route to the destination D is created. If it already exists, it is updated if the sequence number for D in the RREP is larger than the sequence number for D in the routing table, or if the sequence numbers are the same but the new path metric of the RREP is better. The transmitter of the RREP is the next hop in the forward route to D .

If the routing table entry for the forward route has been created or updated, the current mesh point will forward the updated RREP towards the originator of the route discovery. The originator is indicated in the source address field of the RREP message. If the current mesh point is the originator, it can now start forwarding the buffered data frames destined to D .

4.1.5 Optional Maintenance RREQs

Due to the dynamic nature of the wireless medium in a mesh network, an established path, which had the best radio-aware path metric during route discovery, may become worse or alternative paths between source and destination with better path metrics may become available. In order to maintain a best metric path between nodes or to switch to another path with a better metric, HWMP defines the optional implementation feature of so-called *maintenance route requests*.

An active source node with this feature sends a RREQ message periodically for the destinations it is communicating with and for which the path metric has not been updated for a certain time. Since a maintenance RREQ has to discover a current path metric, only the destination should reply to these route request messages. Therefore, the destination only flag is always set in maintenance RREQs ($DO=1$).

Although it is possible to send a separate maintenance RREQ to every destination, the capability of HWMP for multiple simultaneous destinations in a single RREQ decreases the routing overhead.

Maintenance RREQs are processed according to the general processing rules for route request messages.

4.1.6 Route Errors

Links between two mesh points can break, especially in such a challenging environment as the wireless transmission media. HWMP uses route error messages in order to inform all affected mesh points of link breaks, even if they are several wireless hops away.

When a mesh point N detects a link break to one of its direct neighbors, say mesh point M , it will generate a *route error message (RERR)* and sends it to all neighboring mesh points that have paths through N with M as next hop. This RERR message together with the subsequent forwarding of updates of this RERR by the recipients will inform all source mesh points with active paths over the broken link $N-M$ of the link break.

The structure of a RERR information element is shown in Figure 9. The list of destinations contains all mesh points to which mesh point N has an active routing table entry and which now cannot be reached due to the link break. That is, all destinations with M as next hop.




Figure 9. Structure of HWMP route error (RERR) information element

Before the RERR message is transmitted, the destination sequence numbers of the affected routing table entries are incremented by 1, the entries are marked as invalid, and the lifetime field is updated and reinterpreted as the time this invalid entry should not be deleted.

Another mesh point Q that receives a route error message generates and sends a RERR message if necessary, updates its routing table accordingly, and initiates a new route discovery if it is a source mesh point of a broken path.

The list of destinations in the new RERR will be a sublist of the destinations in the received RERR. It contains only those entries of the received list for which the mesh point Q has a valid routing table entry with the transmitter of the RERR as next hop. The destination sequence numbers are taken from the received RERR message. The routing table entries of the affected destinations are invalidated and the corresponding lifetimes are updated in the same way as described above.

4.2 Proactive Extensions of HWMP

In some anticipated usage scenarios, a large proportion of the traffic will be destined for only one or only a few mesh points. For instance, most of the traffic will be destined to one or several mesh portals in a wireless mesh network that provides access to a wired infrastructure and the Internet. Proactive routing to the mesh portals is useful in this kind of usage scenarios.

Mesh points, usually mesh portals, can be configured to periodically broadcast mesh portal announcements through the wireless mesh network. A tree with the mesh portal as root node is build with the same distance vector methodology as used in RM-AODV. Furthermore, messages of RM-AODV are reused for the proactive extension where possible.

The use of this proactive extension to RM-AODV, the reactive part of HWMP, is configurable per mesh portal. This means that mesh portals of the same IEEE 802.11s wireless mesh network can operate with or without the proactive extension.

In order to use the proactive extension, at least on mesh point, usually a mesh portal, has to be configured to periodically broadcast *mesh portal announcements (RANNs)*. This triggers a root

selection and arbitration process, out of which a single *root mesh portal* evolves. The decisions in this process are based on configured priorities, where 0 is the highest one. The MAC address is used as tie-breaker. The mesh portal with the smaller MAC address becomes the root portal. The root portal can be thought of “superior” mesh portal. It is the default mesh portal.




Figure 10. Structure of HWMP mesh portal/root portal announcement (RANN) information element

If a mesh portal is configured to use the proactive extension, it will periodically broadcast mesh portal announcement messages. Figure 10 shows the structure of the mesh portal/root portal announcement information element.

The *mesh portal address field* contains the MAC address of the mesh portal that generated this RANN message. The *mesh portal sequence number field* contains its destination sequence number which is taken from the same sequence number counter as the sequence numbers for the reactive routing. Two flags are defined in a RANN message. The *announcement type flag (AN)* distinguishes between announcements of non-root mesh portals (AN=0) and of the root portal (AN=1). Based on this flag, mesh points can recognize the root portal. RANNs with AN=1 are also called *root portal announcements*. The *HWMP registration flag (RE)* distinguishes between two different modes how a RANN is processed in mesh points. Both non-registration mode (RE=0) and registration mode (RE=1) are explained below. The *priority field* contains the configured priority of the mesh portal for the root portal selection and arbitration process.

The *hop count field* and the *metric field* are initialized with 0 at the mesh portal. They contain the hop count and the accumulated radio-aware path metric to this mesh portal as the mesh portal announcement is propagated through the wireless mesh network. The *lifetime field* contains the lifetime of the path to this mesh portal. The *topology maintenance policy field* defines a maintenance policy for the paths to the mesh portal in registration mode.

A RANN can be thought of as a RREQ requesting paths to all mesh points of the wireless mesh network. This sets up a forwarding tree towards the mesh portal by broadcasting the portal announcements. However, the processing of a RANN is quite different to the RREQ/RREP mechanism in the reactive part of HWMP.

The processing of mesh portal/root portal announcements depends on the setting of the HWMP registration flag in the received portal announcement. If the HWMP registration flag is not set (RE=0), the actions for non-registration mode are performed. If the HWMP registration flag is set (RE=1), the actions for registration mode are performed.

4.2.1 Non-registration Mode

The intention of the non-registration mode is a “lightweight” HWMP topology formation where the routing overhead for the proactive extension is kept at a minimum. The broadcast RANN messages set up a tree that contains paths from all mesh points to the announced mesh portal/root portal, but mesh points are not registered proactively at the root portal.

When a mesh point *N* receives a mesh portal/root portal announcement, it creates or updates a routing table entry to the transmitter of the RANN message. The hop count (incremented by 1) and the path metric (addition of link metric from transmitter to *N*) are updated so that they reflect hop count and path metric from *N* to the corresponding mesh portal. If there is no routing table entry to the announced mesh portal, a new entry is created. An existing routing table entry to the mesh portal is only updated for newer or better path information in a similar way as in the processing of RREQs described above. If the routing table entry to the mesh portal has been created or updated, the transmitter will become the next hop and therefore the parent mesh point in the tree to the mesh portal (see also Figure 11).




Figure 11. Setup of proactive tree to mesh portal P with mesh portal/root portal announcements (RANN)

The immediate update of the routing table is the only action in non-registration mode during the processing of a RANN. If the RANN contained newer or better path information to the mesh portal, the updated RANN message is broadcast to all neighbors of the mesh point.

If bi-directional communication is needed between a source mesh point and the root portal operating in non-registration mode, the source mesh point may send a gratuitous RREP before the first data frame in order to register its address with the root portal on demand.

4.2.2 Registration Mode

Mesh points that receive a root portal announcement from the root portal are registered proactively at the root portal in registration mode.

When a mesh point *N* receives a root portal announcement with the HWMP registration flag set (RE=1), it buffers the RANN and waits for a pre-defined period for other RANNs (of the same announcement round) to arrive. After the expiry of the period, the

mesh point may send out a broadcast RREQ with the destination only flag not set (DO=0) and with the TTL optionally set to 1. This RREQ is requesting a path to the root portal. The mesh point has learned the MAC address of the root portal from the RANN. The RREQ is used for revalidating the paths to the possible parents that have been learned through the reception of the root portal announcements from the neighboring mesh points. The mesh point N chooses the neighboring mesh point with the best path metric to the root portal as its parent in the tree.

The mesh point N now registers itself and its associated IEEE 802.11 stations with the root portal by sending a gratuitous RREP to the root portal. The mesh point sets the first source address field in the RREP to its own MAC address and the following source address fields with the MAC addresses of its associated IEEE 802.11 stations.

After the completion of the successful registration with the root portal, the RANN message that has been chosen for the path to the root portal is updated. The hop count is incremented by 1 and the metric of the link to the parent mesh point is added to the metric field. The updated RANN message is broadcast to all neighbors of the mesh point.

Topology maintenance uses directed, unicast RREQs, which are sent periodically, in registration mode. Optimization of paths to the root portal uses broadcast RREQs with TTL=1 which are sent periodically. The optimization interval is usually larger than the maintenance interval. HWMP defines four maintenance policies that correspond to different tradeoffs between the optimality of the paths to the root portal and the frequency and amount of maintenance updates, i.e. the maintenance overhead.

4.2.3 Hybrid Routing

The so-called *hybrid routing* can only occur when a root portal has been configured (root portal announcements with AN=1 are broadcast) and registration mode is used.

When a mesh point S wants to send data to mesh point D but has no path to D in its routing table, S may send the data frames to the root portal immediately instead of initiating a route discovery for D . The root portal recognizes that D is inside the mesh network, since it knows all mesh points due to the registration mode. It forwards the data frame to the destination together with an indication, that both S and D are “intra-mesh.” This triggers a route discovery for S in mesh point D . This will setup the optimal path between mesh points S and D , on which the subsequent data frames will be forwarded.

Compared to a completely reactive route discovery procedure, there is no latency for the first data frames of an intra-mesh communication with this hybrid routing scheme. The (usually non-optimal) path up and down the tree to the root portal can be used during the setup of the optimal path.

4.3 Support of IEEE 802.11 Stations

IEEE 802.11s WLAN mesh networks with HWMP will support conventional IEEE 802.11 WLAN stations. Mesh APs will generate and manage routing messages on behalf of the WLAN stations that are associated with them. The sequence numbers for the associated WLAN stations are maintained by the corresponding mesh AP. The mesh AP will initiate the route discovery if it does not

have a path to the destination of the data coming from the WLAN station. The mesh AP will also answer RREQs asking for a route to an associated WLAN station with a RREP. The mesh APs will also register their associated IEEE 802.11 stations with the root portal if the proactive extension is used and registration mode is configured.

5. RADIO-AWARE OPTIMIZED LINK STATE ROUTING

The *Radio-Aware Optimized Link State Routing protocol (RA-OLSR)* is an optional proactive routing protocol of the emerging IEEE 802.11s framework. It is an adaptation of the well-known *Optimized Link State Routing Protocol (OLSR)* [10] to the IEEE 802.11s environment. RA-OLSR follows closely the original specification of OLSR [6]. However, there are some differences. RA-OLSR uses MAC addresses instead of IP addresses. The shortest path algorithm uses a radio-aware metric instead of the hop count metric. Therefore, a metric field is added to all topology information messages. Routing support for IEEE 802.11 stations associated to mesh APs is defined. Moreover, a frequency control for link state flooding can be utilized.

5.1 Multi-Point Relays

The major problem of proactive link state routing protocols for wireless mesh networks is the routing overhead caused by the necessary (proactive) network-wide distribution of link state information. The wireless links have only limited capacity. Moreover, they change more frequently than wired links which increases the number of link state updates.

An optimized broadcast mechanism is the core concept of RA-OLSR. Each mesh point selects so-called *Multi-Point Relays (MPRs)* among its direct neighbors, so that every 2-hop neighbor of a mesh point will receive broadcast messages even if only the MPRs forward the broadcast message. This can reduce the number of broadcast messages.

Each mesh point periodically broadcasts *hello messages* that are not forwarded (TTL=1). Hello messages contain a list of the neighbors of the sender. This information allows each mesh node to learn its 2-hop neighborhood. Moreover, the bi-directionality of the links can be verified. The status (asymmetric, symmetric) is attached to each link. Each mesh point also announces its willingness to forward packets in hello messages. The willingness is in the range from “will never” to “will always”. RA-OLSR stores the local neighborhood information in several information repositories, which are the *link set*, the *neighbor set*, and the *2-hop neighbor set*.

Each mesh point selects its MPRs independently and solely based on the received information about its 2-hop neighborhood. The only requirements are that the complete 2-hop neighborhood receives broadcast messages if only MPRs forward them, and that only neighbors with symmetric links and willingness to forward are considered.

RA-OLSR proposes a simple heuristic for the MPR selection [1]. This is the adaptation of the original proposed OLSR MPR selection procedure [6] to IEEE 802.11s. Other MPR selection algorithms are also possible.

The selected multi-point relays are stored in the *MPR set*. The MPR set has not to be minimal, but the smaller the MPR set the lower the overhead. Neighbors that have been selected as MPR by a mesh point will have a link status indicating the MPR selection in the hello messages. A mesh point can derive from this information the mesh points that selected it as an MPR. These mesh points are called *MPR selectors* and are stored in the *MPR selector set*.

5.2 Forwarding of Broadcast Messages

All broadcast messages are forwarded throughout the wireless mesh network according to the *default forwarding algorithm*. The default forwarding algorithm ensures that only MPRs forward the broadcast message, that every MPR forwards the broadcast message only once, and that only broadcast messages with large enough TTL are forwarded.

5.3 Topology Information Dissemination

Each mesh point that has been selected as MPR periodically broadcasts *topology control (TC) messages* in order to distribute its link state information within the wireless mesh network. The TC message contains a list of neighbors of the originating node. It must contain at least all MPR selectors of this node. An *advertised Neighbor Sequence Number* is associated with the neighbor list. This allows to recognize out-dated topology information. TC messages are forwarded in the mesh network according to the default forwarding algorithm.

A further optimization of the dissemination of topology control messages is a frequency control for broadcasting link state information as known from Fisheye State Routing [9]. Nearer mesh points receive topology information more often than further away nodes. In order to achieve this, the TTL in subsequent TC messages alternates between 2, 4, and maximum TTL.

A mesh point stores the information of received TC messages in the *topology set*.

A classical shortest path algorithm computes the entries of the RA-OLSR routing table from the link set, neighbor set, 2-hop neighbor set, and topology set based on the radio-aware link metric. The routing table contains entries for all reachable destinations in the mesh network since RA-OLSR is a proactive routing protocol. The routing table has to be recomputed if any of the above information sets has changed. Furthermore, it might be useful to propagate these changes immediately by sending a hello or TC message.

All topology information that is stored in the repositories, such as, the topology set, has an expiration time associated with it in order to provide some robustness against the loss of RA-OLSR routing control messages.

5.4 Support of IEEE 802.11 Stations

Each mesh access point maintains a *local association base (LAB)*. It contains a list of all IEEE 802.11 stations that are associated with this mesh AP. The mesh AP distributes its association information in the mesh network by periodically broadcasting *local association base advertisement (LABA) messages*, which are forwarded according to the default forwarding algorithm. Each mesh point stores the association information of the received LABA

messages in its *global association base (GAB)*. Both LAB and GAB are used for the computation of routing table entries for IEEE 802.11 stations associated with a mesh AP.

The association bases are organized in blocks of local association tuples. In order to save bandwidth, a *local association base checksum advertisement (LABCA) message* may be sent instead of a LABA message. A LABCA message contains only the checksums of the LAB blocks. If there is a mismatch between a received checksum and the corresponding checksum in the GAB, the mesh point requests an update of this LAB block from the originating mesh point by sending an *association base block request (ABBR) message*. This behavior is called *checksum diffusion mode*, which is optional, in contrast to the *full base diffusion mode*.

6. AIRTIME ROUTING METRIC

The proposed IEEE 802.11s amendment [1] defines a default radio-aware routing metric for basic interoperability between IEEE 802.11s devices. The *airtime link metric* is a measure for the amount of the consumed channel resources when transmitting a frame over a particular wireless link.

Equation (1) is used for the calculation of the *airtime cost* c_a of each link. The path metric is the sum of the metrics of all links on the path.

$$c_a = \left[O_{ca} + O_p + \frac{B_t}{r} \right] \frac{1}{1 - e_{fr}} \quad (1)$$

The *channel access overhead* O_{ca} , the *MAC protocol overhead* O_p , and the *number of bits* B_t *in a test frame* are constants. Their values depend on the used IEEE 802.11 transmission technology such as IEEE 802.11b or IEEE 802.11g. The *transmission bit rate* r in Mbit/s is the rate at which the mesh point would transmit a frame of size B_t with *frame error rate* e_{fr} , based on the current conditions of the radio environment.

7. EXTENSIBILITY

In order to have the flexibility to choose an optimal path selection protocol for the anticipated usage scenario, but still to have some degree of interoperability between different vendors, the proposed IEEE 802.11s amendment defines an extensibility framework.

Beacons of mesh points contain a *path selection protocol identifier* and a *metric identifier*. They indicate the active path selection protocol and the active routing metric currently used in an IEEE 802.11s WLAN mesh network. A mesh point that wants to join an existing IEEE 802.11s WLAN mesh network has to be able to support the announced path selection protocol together with the announced routing metric. If not, it cannot join the mesh network.

Interoperability is achieved by the requirement that every IEEE 802.11s compliant device has to implement the default routing protocol of the IEEE 802.11s amendment, HWMP, as well as the default routing metric, the airtime link metric. This means, that every mesh point, no matter who the vendor is, can “speak” HWMP and the airtime link metric. IEEE 802.11s WLAN networks with this configuration will provide complete interoperability.

bility. Interoperability is further increased by the broad applicability of HWMP to many usage scenarios.

A path selection protocol identifier is four octets long and consists of an *organizational unique identifier (OUI) field* and a *path selection protocol identifier field*. The OUI allows the use of vendor specific path selection protocols.

Only one path selection protocol can be active in an IEEE 802.11s WLAN mesh network at a time.

The extensibility framework with this mechanism allows the use of other routing protocols and/or routing metrics that are better suited for some scenarios instead of the default ones. Figure 12 illustrates the extensibility with respect to path selection protocols.




Figure 12. Extensibility of IEEE 802.11s with respect to path selection protocols

8. OUTLOOK

The paper presented a detailed overview on the proposed routing of the upcoming IEEE 802.11s standard for WLAN mesh networks. The configurability of the default routing protocol HWMP and the extensibility framework for the routing with RA-OLSR as optional standardized routing protocol and the ability to integrate optimized and vendor-specific routing protocols gives IEEE 802.11s a broad applicability to many usage scenarios of wireless networks.

The presented material is based on the very first draft standard of IEEE802.11s [1] which will change until it is finally approved. Nevertheless, the general concepts for the routing framework and for HWMP and RA-OLSR are agreed on and are quite stable. This justifies a publication like this, even if it is very likely that details will change.

The task group “s” is actively reviewing and continuously improving the draft standard. Contributions have been announced in

response to comments from a first internal review. Many comments and improvements are expected during the first letter ballot later this year. The final approval of the IEEE 802.11s standard is expected for 2008.

9. ACKNOWLEDGMENT

I thank my wife Claudia for the support during the preparation of this paper. I thank Christian Bettstetter who organized this conference session. Furthermore, I thank my Siemens colleagues of the ad hoc mesh team – Michael Finkensteller, Matthias Kutschenreuter, Rainer Sauerwein, Christian Schwingenschlögl, and Norbert Vicari. Last but not least, I thank my “IEEE 802.11s colleagues.”

10. REFERENCES

- [1] IEEE P802.11s™/D0.01, Draft amendment to standard IEEE 802.11™: ESS Mesh Networking. *IEEE*, March 2006, work in progress.
- [2] IEEE Wireless LAN Edition, A compilation based on IEEE Std 802.11™-1999 (R2003) and its amendments. *IEEE*, 2003.
- [3] Akyildiz, I. F., Wang, X., and Wang, W. Wireless mesh networks: a survey. *Computer Networks*, vol. 47, no. 4, March 2005.
- [4] Aoki, H. et al. 802.11 TGs Simple Efficient Extensible Mesh (SEE-Mesh) Proposal. *IEEE P802.11 Wireless LANs*, Document IEEE 802.11-05/0562r0, June 2005.
- [5] Bahr, M., Wang, J., and Jia, X. Routing in Wireless Mesh Networks. In *Wireless Mesh Networking: Architectures, Protocols and Standards*, Zhang, Y., Luo, J., Hu, H., Eds., Auerbach, Dec. 2006, to be published.
- [6] Clausen, T. H. and Jacquet P., eds., Optimized Link State Routing Protocol (OLSR). *IETF Experimental RFC 3626*, Oct. 2003.
- [7] Conner, W. S. IEEE 802.11 TGs Usage Models. *IEEE P802.11 Wireless LANs*, Document IEEE 802.11-04/0662r16, Jan. 2005.
- [8] Hauser, J., Baker, D., and Conner, W. S. Draft PAR for IEEE 802.11 ESS Mesh. *IEEE P802.11 Wireless LANs*, Document 11-04/0054r2, Jan. 2004.
- [9] Iwata, A., Chiang, C.-C., Pei, G., Gerla, M., and Chen, T.-W. Scalable Routing Strategies for Ad hoc Wireless Networks. *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 8, Aug. 1999.
- [10] Jacquet P. et al., Optimized Link State Routing Protocol for Ad Hoc Networks. In *Proc. INMIC 2001*, 2001.
- [11] Johnson, D. B., Maltz, D. A., and Hu, Y.-C. The Dynamic Source Routing Protocol for Mobile Ad Hoc Networks (DSR). *IETF Internet Draft*, draft-ietf-manet-dsr-10.txt, July 2004, work in progress.
- [12] Perkins, C. E., Belding-Royer, E. M., and Das, S. R. Ad hoc On-Demand Distance Vector (AODV) Routing. *IETF Experimental RFC 3561*, July 2003.