IEEE P802.11
Wireless Access Method and Physical Layer Specifications

Title: Required MAC functions to support multirate PHY's

Authors:

Wim Diepstraten
AT&T GIS (NCR)
WCND-Utrecht
Nieuwegein The Netherlands
Tel: (31)-3402-97482
Fax: (31)-3402-97555
Email: Wim.Diepstraten@utrecht.ncr.com

Barry Dobyns
Photonics.
2940 North First Street
San Jose, CA 95134
Tel: (408) 955-7930
Fax: (408) 955-9750
Email: dobyns@acm.org

Pablo Brenner
LANNAIR Ltd.
Atidim Technological Park
Devorah Haneviah St - Neve Sharet
Tel Aviv 61131 - ISRAEL
Voice: +972-3-6459127
FAX: +972-3-5447146
EMail: pablo@lannet.com

Rui Valadas
University of Aveiro
Dept. Electronics and Telecom.
3800 AVEIRO
PORTUGAL
Tel: +351 34 381937
Fax: +351 34 381941
Email: rv@ua.pt

Abstract:
This submission contains a discussion of the functions that are needed in the MAC to support MultiRate PHYs, the possible solutions, and a proposal for the minimum set of hooks in order to support this option.

Related Documents:

IEEE DOC P802.11-94/119: Gear Shifting Proposal
IEEE DOC P802.11-93/157: Proposed Changes to Draft Standard in order to support MultiRate PHYs

Why Multiple-Rates

There are several reasons for providing this functionality:

- Extensibility
- Speed/Range tradeoff
- Speed/Power Consumption tradeoff

Extensibility

The extensibility requirement is based on the fact that the Wireless technology is advancing rapidly, and higher rates are foreseen in the future. We already see this in the current standard, where in the Frequency Hopping PHY, which while the specification for a 1 MBit/s PHY is not yet finished, work is already being done for a 2 MBit/s FH PHY.

So higher rates are foreseen in the not so far away future, and a higher rates migration path should be provided, especially in a wireless environment, where the bandwidth needs to be shared between any existing equipment following an existing standard, and equipment that implement a newer (faster) standard.

While there is almost no argument on the need for extensibility, there are lots of discussions whether such extensibilty must be provided by a Different Rate ESS, a Multi-Rate BSS or by a Multi-Rate ESS (hence adding a new Overlapping BSS with higher rate capabilities).

Fixed Rate per ESS approach:

This is a very inflexible migration path, which forces a whole installation to switch to a higher speed, only when all stations are capable of doing so. This will force users to make their old equipment obsolete when they want to go to a higher speed. It does further not allow ESS overlap on the same channel, because that would yield problems in the NAV calculation in the overlap area.

We do not recommend this approach.

Fixed Rate per BSS approach:

This approach indicates that each BSS has a fixed rate and if two rates are to be supported in the same geographic area, it will be done by two overlapping BSSs, each operating at a different speed. This configuration requires that the PHY's should support multiple rates, for the Roaming function, while the roaming algorithm will need to cope with it. This configuration has the same NAV calculation problem as any mixed rate overlap environment (the low speed BSS stations will not understand the high speed NAV duration fields).

The authors believe that most of the problems indicated for the Multi-BSS approach, are the same for the Multirate BSS approach, and relates to the NAV update issue.

In terms of medium efficiency this approach is also less desirable, because a number of frames need to be transmitted at both rates so in each of the overlapping BSS's. Examples of this are all broadcast and multicast frames, and Control Frames such as Beacons.

We do not recommend this approach.

MultiRate BSS approach:

The rational behind this approach is to let different bitrate stations coexist on the same BSS. This requires that all control and Broadcast/multicast frames need to be transmitted in a pre-agreed rate (supported by all stations in the BSS). Higher rates can be used for unicast data frames to those stations that support the higher rate.

The proposal is flexible enough to allow for "high-speed only" BSSs or low rate only BSSs.

It should be noted that by allowing "higher rate capable" stations to transmit in higher rates, we are also improving the performance of the "lower rate only" stations, this is possible because the "higher rate" stations release the medium faster, hence leaving the medium to the other stations.

We recommend to adopt this approach.

Speed/Range Tradeoff

Experience shows that sometimes, the topology limitations are such that even while most of the stations can work properly on a high bit rate, some of them may (even temporarily) have to switch to lower rates to keep a given Bit Error Rate level.

Mapping in advance the user environment to locate the APs such that ALL stations will be able to run ALWAYS in the higher speed rates will be almost impossible, and the area will have "black spots" or the user will have to put more APs than required.

A MultiRate BSS could accomodate this kind of situations, while a Multi BSS approach cannot.

Speed/Power-Consumption Tradeoff

In the current IR PHY there is no speed/range tradeoff but there is a speed/consumption tradeoff.

The same range is being specified for both data rates with the purpose of minimizing the hidden station problem. the immediate consequence is that the higher data rate requires more power.

Most stations may not want to pay the price of a higher consumption., probably in the near term the higher rate will mainly be used by APs. In that case an asymmetric behavior will result in that traffic from the AP to stations will be at the higher rate and traffic from the stations to the AP at the lower rate.

It is also foreseen that with the advances in the optoelectronic technology more efficient components will be developed in the long term allowing for a migration of all transceivers to the higher rate.

MAC Concerns:
In discussions about the functionality of a multi-rate MAC, the following subjects must be addressed and analysed for proper operation.

1. Control Packets (Probes, Beacons, etc) must be heard by all stations
2. Multicast/Broadcast must be heard by all station
3. How does a station know which rate to use for transmission?
4. How do stations update their NAV for transactions in other speeds? What is the effect of fragmentation.
5. How are hidden-stations affected?
6. How does CCA work for different speeds ?
7. How are Time Bounded Services Affected ?
8. What hooks are needed to allow a flexible multi-rate solution?

The above subjects are analysed, to determine what the basic requirements are for a flexible multi-rate implementation. The following set of assumptions are used, to clarify the discussions.

Definitions

BSS_BASIC_RATE_SET:

A set of rates that all the stations on the given BSS are required to be capable of reception. According to the PHYs definitions the default BSS BASIC RATE SETs for the different PHYs will be:

For DS: {1,2}
For FH: {1}
For IR: {1,2}

Note that this value is preset for all stations in the BSS

STATION_BASIC_RATE:

A value belonging to the BSS BASIC RATE SET, that is used by the station for specific transmissions (it could change dynamically, for example the Station Basic Rate on the IR depends on the Power Consumption Mode of the Station).

EXTENDED_RATE_SET:

The set of rates beyond the BASIC_RATE_SET that a station supports. This can be a speed that is defined in future PHY standards.

PLCP_RATE:

This is transparent to the MAC, but must be the same for all stations in the BSS. This is the basic assumption that is already followed in the current multi-rate PHY proposals. It allows stations to automatically detect the speed of an incomming frame, and decode/receive it at the intended rate if it is supported by that PHY.

Other suggested assumptions are:

- The **Preamble** and the **PLCP Headers** are transmitted always at the **PLCP_RATE.** The PLCP Headers contain both dataphase rate and length information. The length information in the PLCP Header should be specified in "Time" to assure that all stations independent of their implemented bitrate can calculate the end of the frame in a bitrate independent way. So if it receives a header of a (higher speed) frame for which it does not understand the dataphase modulation, it is still able to calculate how long the medium will be busy, to assert a proper CCA signal to the MAC.

- The different **IFS**s are the same for the whole ESS.

- All the **Control, Multicast and Broadcast Messages** are sent on the **STATION_BASIC_RATE** (which as specified above belongs to the BSS BASIC_RATE_SET). This solves MAC problems 1 and 2.

- All RTS and CTS frames are transmitted on the STATION_BASIC_RATE. This allows reception of the duration information by all stations.

- The duration field in the RTS/CTS frames are specified in time (usec) rather then bittimes. This and the previous assumption solves MAC problem 4. Remote stations do not need to interpret the bitrate, but just use the value in the "duration" field to update their NAV. The destination station that is returning the CTS, does only need to substract the CTS duration from the "Duration" field under normal sircumstances.

- All unicast data frames are sent on any rate. The algorithm for selecting this rate is implementation specific. It is something that is determined by the transmitter only, which can use any rate that the PHY of that station is capable of using.
  In general three approaches can be followed:
  - Transmitter gathers information on rate supported by the remote, and uses the supported rate for subsequent transfers to that destination.
  - Used rate is negotiated in an RTS/CTS exchange per frame.
  - Transmitter decides the transmission rate based on the own station characteristics (e.g. Power Consumption Mode).

  Again the algorithm for selecting this rate does not need to be standardized.
  Some trivial algorithms could be:
  1. Try high, retransmit on lower (go back to high after T time).
  2. Keep "Supported rate" table for each peer.
  3. Keep dynamic tables for each peer using a signal quality (or any other parameter) dependent algorithm.
  4. Transmit always in BASIC_RATE.

Above, two possible methods are suggested to use multiple rate. The basic idea is that a MAC transmitter is controlling the rate at which it is desired to transmit, depending on its capabilities, while the selected rate should be supported by the intended receiver, over a link that has sufficient quality to achieve sufficient performance.
To allow dynamic rate switching the transmitter would preferably need to know the following information:

- The supported rate of the local PHY.
- The supported rate of the destination station.
- Link quality conditions of the link to the destination.

Although there are ways to gather part of this information without any remote station involvement, it is considered desirable for the standard to have the hooks specified to gather that information.
In an infrastructure network "supported rate" information can easily be exchanged in the association process.
In ad-Hoc networks such information could be gathered by specifying a probe function to gather this sort of information. It could also be that the supported rate capabilities are advertised by an ad-Hoc station in the distributed Beacon.
This can be accomplished in the standard by specifying a **"Supported_Rate** element for this purpose that needs to be included in:

- The "Association Request"
- The "Association Response"
- The "Beacon" (both Ad-Hoc and Infrastructure)

The **"Supported_Rate"** element can be used in a management Request/Response PDU as a probe frame to allow stations to gather the required information to build their database.
The element should specify both the supported transmit aswell as receive rates.

Apart from the above described method that does not need any support of a receiving station, an active rate negotiation procedure can be specified in the standard, this active negotiation is important for the cases when medium conditions may change (e.g. mobile stations)

This requires a procedure in which the receiver is participating in the rate determination. This requires the definition of two additional elements for the exchange of information between transmitter and receiver much in the same way as specified in [2].
The following element definition, or specific field definition is required to support this method:

- **"Requested_Rate"** element or field in the RTS frame (one rate only).
- **"Granted_Rate"** element or field in the CTS frame (one rate only).

- **Data Messages** sent **with RTS/CTS** can go through a rate negotiation procedure by using **"Requested_Rate"** and **"Granted_Rate"** elements in the RTS and CTS frames respectively.

- The RTS frame can contain the "**Requested_Rate"** element, and a duration field that specifies the duration in usec based on the requested datarate.

- The **CTS** can contain the "**Granted_Rate"** element, which indicates the bit rate at which the Data transmission should take place, together with a duration field that specifies the length of the transaction based on the granted rate.
  If the Destination Station does not support the "Request Rate" then the "Granted_Rate" will belong to the BSS_BASIC_RATE_SET.
  The decision about this rate can be derived by a MAC by interpreting a "Signal Quality" status from the PHY.

- **Data Messages** belonging to an RTS/CTS MPDU are always sent at the **Granted_Rate** of the CTS, which is either the Requested_Rate or the BASIC_RATE as specified above.
  This solves a potential problem that would occur if a station heard the CTS but doesn't hear the Data Trasnsmission, hence assuming a shorter NAV.

How do stations update their NAV?

Given that the "Duration" field is given in usec, and sent at the Basic_Rate, a station only has to use this value to update its NAV. Stations have to update the NAV on both the RTS and CTS frames, in which the CTS duration field will show the correct NAV value for the "Granted_Rate". They can do this without the need to know at which speed the data will be send.

What is effect of fragmentation?

When fragmentation occurs, then the Data and Ack frames do contain "Duration" information. Assuming that the Ack is transmitted at the basic rate, then all stations around the intended receiver will hear the relevant "Duration" of the next fragment, as it is copied from the "Duration" field of the previous data fragment. This data fragment however is possibly transmitted at an extended rate, so that not all stations can decode its "Duration" field. The effect of this is that stations around the transmitter, but not in range of the receiver sending the Ack, do not have NAV information. So they are in carrier sense range to hear the fragment, and could start transmission after a DIFS + backoff slot period following this fragment. Depending on the length of the DIFS, this can overlap with the Ack, which can result in Ack failures.
This does not occur when the DIFS = 2*SIFS+Ack.

This effect is basicly the same as the problem when RTS/CTS is not used.

How are hidden stations affected?

This subject is only relevant for the situation where stations can hear the RTS frame but not the CTS frame. When the requested rate is not granted by the destination, then there will be a mismatch between the NAV specified in the RTS frame, and the required NAV (which should become longer). Those stations however will hear the subsequent dataframe, and will defer until at least DIFS after that. This DIFS may not be long enough to prevent that channel access already starts during the Ack which is outside the hearing distance. Ideally those stations should at least defer until 2*SIFS+Ack duration. Conclusion here is that there migth be a potential Ack overlap problem, depending on the DIFS parameter. This issue is resolved when DIFS >= then 2*SIFS+Ack. Again this issue is independent of multi-rate support.

How does CCA work at different speeds?

The current DS and FH PHY's do specify a PLCP header that contain "Current_Rate" and "Length" information. This allows them to use different speeds in the Dataphase, while all headers can be received by all receivers independent on their rate capabilities.

The current "Length" field does however likely specify octets, or bits. To make the interpretation of the length field bitrate independent it is required that this field is specified in "Time" (usec) rather then bits or Bytes.
The importance of this is the following functionality that should be enforced in the standard:

- One of the functions of the "Length" field should be that it determines when the CCA is indicating a free medium. So when a station receives a PLCP header of a frame with a future x Mbps, then the length field will determine when the CCA is raised. This prevents that for instance CS detection mechanisms fail because they are not able to detect the modulation used in the dataphase of that frame.
- The other function is to use it as the end delimiter in the intended receiver, to subsequently signal to the MAC where the frame ends, so where the CRC is located.

With the above PHY functionality, coexistance can be assured in a future environment, with extended rates such that the CCA will work independent of the bitrate.

Summary:

To support multiple rate within a BSS in a flexible way, a basic set of rules are defined together with 3 element codes.

- **"Supported_Rate"** element
- "**Requested_Rate"** element
- "**Granted_Rate"** element

The "**Supported_Rate"** should be included in every Associate Request, Associate Response and Beacon. It can also be used to gather the supported rate information using the Request/Response management frames.
The above defined rules provide a migration path to future higher rate technologies, while introducing little overhead in the current MAC and PHY specifications.

References

[1] IEEE DOC P802.11-94/119: Gear Shifting Proposal, P. Brenner, M. Rothenberg - LANNAIR

[2] IEEE DOC P802.11-93/157: Proposed Changes to Draft Standard in order to support MultiRate PHYs, P. Brenner - LANNAIR