IEEE P802.11

Wireless Access Method and Physical Layer Specification

MAC Definition Conflicts, Gaps and Other Issues

Michael Fischer
Digital Ocean, Inc.
4242–3 Medical Drive
San Antonio, TX 78229
Telephone: +1–210–614–4096
Facsimile: +1–210–614–8192
email: mfischer@CHILD.com

Abstract

This document contains a categorized list of ambiguities, contradictory provisions, missing information, and other open issues on the 802.11 MAC protocol. These descrepancies and other issues were encountered by the author while working on the draft formal description of the MAC.

Introduction

In the process of generating the draft formal description of the 802.11 MAC protocol, the author encountered numerous descrepancies in the prose description of the MAC. (In particular, in the D5.1 updates as of the end of the November, 1996 plenary meeting.) Some of these discrepancies are simply artifacts resulting from incomplete insertion of approved changes, or incomplete removal of obsolete terminology or features. However, others are truly ambiguous as written, due either to conflicting statements or missing information. If not resolved, these problems in the text could, and probably will, contribute to improper and/or non-interoperable implementations. While the formal description, being added as normative Annex C, will {necessarily} clarify most of these issues, it is important for the text and formal descriptions to match as closely as practical. Also, in some cases a "judgement call" was necessary to be able to generate any formal description, and these need to be verified as whether they describe properly the desired MAC behavior.

1. Inter-Function Conflicts

1.1. The existing text says that the backoff does not decrement during the ATIM window, which means that a station with a non-zero backoff count at TBTT would not be able to participate in the ATIM window. This is clearly wrong. In discussions with several people involved in IBSS power management ballot comments, it appears that the intent was to have two sets of backoff and current CW values, one for normal DCF operation and one for ATIM window operation, with the ATIM window values swapped for the DCF values at TBTT (after cancelling a possible pending non-ATIM transmission attempt in progress) and swapped back at the end of the ATIM window. This is how the state machines are currently implemented. Open issues include what to do with the CW values at the boundaries. Currently the state machines treat the two as coroutines, swapping the actual values, which is the simplest (fewest special case tests) and appears to preserve fairness better than if some stations reset their CWs to CWmin while others do not.

1.2. Scanning requires changing the PHY channel selection according to a time and channel list independent from anything to do with BSS timing. Since scanning is permitted when a station is a member of a BSS or IBSS, this implies that there is a facility which permits the scan function to suspend normal DCF activities — otherwise the station might send an ACK on a different channel than the frame reception, might send a frame on a different channel than the rest of the BSS was using, etc. At least, the scan needs to commence when there is no frame exchange sequence in progress, and must prevent any new frame exchange sequences from being started. In an IBSS there is also a need to coordinate with the synchronization function, since permitting a scan to begin during a beacon interval when this station sent the beacon would leave no station to respond to incoming probe requests. The text does not even imply that these conflicts exist.

1.3. If a transmission attempt is suspended due to TBTT in an IBSS or due to scanning, and aMaxMsduTransmitLifetime is exceeded before that suspension ends, what happens to the CW ?

1.4. Nothing prohibits a station from initiating an authentication sequence when a previous sequence is still in progress, however, multiple outstanding shared key sequences imply a cache of challenge text per DA — this is rather a mess. Either this should be prohibited or the need for such a cache should be clearly identified, since using the same challenge text for multiple, concurrent requests destroys the value of the authentication mechanism entirely unless bilateral keys are used, in which case doing so still gives away too much information about the keys in the case of successful authentication.

1.5. Section 8.3.2 says the MIB prevents WepKeyOn=true if open system authentication used, but the privacy MIB may be set before authentication occurs, so the place this needs to be enforced is in the transmit setup encryption routine when selecting the key, and the corresponding place in the receive path.

2. Inadequate Information

2.1. MlmeDeauthenticate.indicate should probably be issued for all causes of unsuccessful authentication, including cases where the station was already authenticated but the new authenticate attempt failed.

2.2. There needs to be a way for SME to specify the capability bits for start, scan, etc.

2.3. What has happened to aBssBasicRateSet — it is gone from the MIB but not replaced?

2.4. There is no way that SME can wake up a sleeping MAC other than by issuing MlmeReset.request, which as severe side effects. Perhaps the time when MlmePowerMgt.request becomes effective should be specified to be immediate, at least in the case of wake-up. Also, since power management state changes can only be announced to the BSS by acknowledged frame exchange sequences, if there is no traffic to/from LLC when power management state is changed, the delay can be arbitrarily long, which is certainly undesirable when changing from active to power save, so there may be a need to generate dummy traffic to carry the announcement.

2.5. The key for use encrypting the challenge text is not clearly identified. The current approach uses the current default WEP key.

2.6. The previous BssId is no longer in the reassociate request. If this is intentional the text should indicate where this information is now obtained.

2.7. What replaces aProbeDelay, which is used in several places in the MAC? What replaces aProbeResponse for management frame response timing?

2.8. Do management frames which are ACKed by for which no response is received get retried? Presumably not, of this would be stated, but a clear statement that this is not the case should be included or the retry strategy should be listed.

2.9. There are quite a few causes of transmit failure which can be identified, but only "undeliverable" to report back to LLC if the MaUnitdata.request was valid when accepted. Also, "provided service class" on reception is ambiguous if some fragments were delivered in the contention period and some in the CF period. Currently the delivery of the last fragment is what is reported updward.

2.10. We treat contention free delivery as a "priority" for LLC, which is historic from the days of TBS, but there is nothing by which an AP can tell whether something sent in response to a CF-poll is of CF-priority and therefore needs to be sent in the CF period. The matter gets even more complex if an MSDU has some fragments delivered in each period.

2.11. Should the NAV duration that exists when a scan request changes from the original channel be counted during the scan, in case the scan duration is short enough that the NAV would still be set when the scan completes?

2.12. The success/failed status for the Mlme primitives does not provide a way to report reason or status codes back to SME. This appears to render those codes useless since they cannot be reported.

3. Redundant Information

3.1. There are no decisions/actions dependent on the station long and short retry counters, only on the per-MSDU versions, therefore the station counters do not appear in the formal description.

3.2. The tuple cache for receive duplicate filtering is not quite a cache in the classic sense. A new entry is created only when a new (address, sequence) pair needs to be recorded, when a new fragment number for an already-cached entry is detected, the existing entry is updated rather than a redundant entry being created.

3.3. aWepDefault=0 should probably go away.

3.4. MaUnitdata.Indication includes a failure status but nothing which causes such to be reported.

3. MIB Issues

3.1. A number of MIB entries in the new ASN.1 version are wrong — especially the WEP default key entries in the privacy group. The full set of issues identified are in the comments of the "32xx" set of SDL pages.

3.2. WEP keys need to be write-only.

3.3. aStationId appears useless as defined, aPrefMaxMpduFragmentLength is a very poor name, aAuthenticationType is of questionable value as defined.

3.4. There are a number of PHY parameters in nanoseconds or milliseconds which should be changed to Usec or Kusec unless there are internal uses in the PHY which require the other units — they are used in the MAC in equations whose results are in Usec or Kusec, so all they do from a MAC point of view is to require rounding.

3.2. The exact conditions under which counters are incremented is inadequately specified. The comments in the counters group in the "32xx" SDL pages specify the exact conditions under which the counters are currently being upated — this should be checked for correctness, then the text transferred elsewhere.

3.3. There appear to be items missing from the MIB which are unnecessary for operation, but may be important for management — especially operational capabilities such as cfPollable and supported data rates.

4. MAC/PHY Interface Issues

4. 1. aMpduDurationFactor is defined Integer32, but for the FH PHY the default value is 1.03125. Since SNMP cannot support Real, and this value will always be in the range 1:2, I suggest using an encoding as an Integer32 containing the fractional part scaled by 1e9, which should be adequate up to about 100Mbps data rates.

4.2. The RTS and ACK timeouts, as well as the transmit rate for CTS and ACK, require knowledge of the receive data rate. The PHY must know this rate, since the signal would not have been demodulated if the rate could not be determined, but this information is not passed to the MAC. I suggest that this should be part of the RxVector of PhyRxStart.indicate, and be encoded in a uniform manner for all PHYs.

4.3. Active scanning is specified to delay for no "medium activity" before sending the probe request. This is currently taken to mean PhyRxStart.Indicate, mainly for simplicity, since the RTS and ACK timeouts work this way. For probe timeout this might want to be PhyCca.Indicate(busy).

4.4. aPreambleLength and aPlcpHdrLength are in units of bits, but the IR PHY has a non-integral number of bits in its PLCP header due to its rate field, and a variable, possibly non-integral number of bits in its preamble. This makes it impossible to calculate the response timeouts for an IR PHY.

5. Observations

5.1. A local LLC entity at an AP would need knowledge of BSS membership to be able to engage in station-to-station transfers. To avoid creating a messy special case for this (uncommon) situation, the current approach is to have MAC Data Service at the AP send everything to Distribution Service, thereby avoiding the need for the AP to need an extra path for transmissions that are not FromDs.

5.2. If supporting multiple outstanding MSDUs, it is necessary to fragment before queueing power save Mpdus for transmission because the time to transmit must be available to select something which will fit in the remaining time of this beacon interval.

5.3. Wep is transparent to LLC, but not to SME because of the ICV failure counter.

5.4. The “shall” in the description of aMaxReceiveLifetime is not achievable.

5.5. The possibility of BSSID conflict in an IBSS raises the issue of what to do if the BssId matches but the SSID does not — there appears to be a way to detect and recover from this occurrence, which should be discussed for possible inclusion. In any case, the detection should be explicitly identified and an error handling if not recovery technique should be listed.