

Tentative MAC Minutes Monday, August 28, 1995

The meeting was called to order by chairman Dave Bagby at 1:25 PM. Carolyn Heide secretary.

Administration: July minutes approved by consensus.

Papers: 95/137, 95/182, 95/183, 95/187, 95/188, 95/189, 95/190

Goals

- process D2.0 LB comments (if any...)
- complete work on clause 8.
- outline conformance statements to correspond to D2.0, flesh out to match D2.0

D2 Letter Ballot Subjects, list of things to discuss this week

These topics will be dealt with this week. The goal is to create papers which can be used by many people as the resolution to letter ballots. Everyone can say "I object to xxxx, and it can be fixed by incorporating the text in document 95/yyy."

Topics:

1. ad hoc power management (small group brings resolution, 20 minutes).
2. conformance statements, what's required? (Vic to provide guidance)
3. random backoff time, subclause 6.2.5. What values? (20 minutes)
4. connection services incomplete: group consensus is that these are editorial things that were not removed when July meeting removed connection services. (editorial)
5. status codes, fields are there and a list could be developed (small group)
6. MIB variables are out of data with the rest of the draft. Clause 8 & 5.3.2. (small group)
7. fragmentation vs. dwell boundaries (summarize in MAC & ask in full working group)
8. state machines point out errors that are in them. Hard to approve when wrong. Previous decision was text overrides state machines when they disagree. That should be stated somewhere also, maybe state machines which should an annex. (group drafts Michael, Simon, Johnny)
9. multirate description inadequate, needs fixing or removing. (20 minutes to identify the problem and ask for help from multirate advocates)
10. subclause 5.2.3 encryption by fragmentation doesn't reflect MPDU vs. MSDU decision, and is that right decision. (20 minutes)
11. use of the TO_DS bit, no definition of how a station makes that decision. Something needs to be said about this. (20 minutes, two straw men)
12. method of generated a TIM with a compressed bit map, compressing the bit map is more complicated than necessary. (7 minutes each side, 10 minute discussion)
13. how is beacon interval time coordinated in an ad hoc network. (10 minutes)
14. relation between reallocation and authentication. need reassociation message from state 2 (10 minutes)
15. combining association and authentication. does it work now, would it be limiting in the future. (5 minutes each opinion, 15 minute discussion)
16. what to do if station leaves, gets disassociation, then returns and it thinks it still owns the SID which has been reassigned. Send disassociation when get SID/SA not match. (20 minutes)
17. ICV algorithm discussion. (20 minutes)
18. TSF timer algorithm is screwed up. (refer to clause 8 group for discussion)
19. how to facilitate key change (20 minutes)
20. should simultaneous MSDU limit be different for AP vs. non-AP, and what happens if maximum exceeded? (20 minutes)

21. beacon response to poll vs. power save mode (20 minutes)

11. use of the TO_DS bit

Proposed Subclause 4.1.2.1.3 replacement: "If a station is operating in an IBSS network the TO_DS bit shall be zero. When operating in an ESS the following rules apply: 1. all management frames shall have the TO_DS bit set to zero; 2. all control frames shall have the TO_DS bit set to zero; 3. when in the associated state, data frames shall have the TO_DS bit set to one, unless the station has explicit knowledge that the destination station is in the same BSS. ~~This explicit knowledge can be gained by getting an RTS from a station known not to be an AP.~~"

No conclusion is reached as to whether or not this replacement text is correct. How this explicit knowledge might be gained is discussed. There is feeling that if it comes from outside the MAC (e.g. human intervention), a MIB variable should be used for conveying it. There is a question as to whether doing nothing to the text which is there hurts anything. Most people feel if the criteria with which this bit is to be set or not is outside the scope of the MAC, the standard should say that.

To be continued ...

Meeting adjourned: 4:30 PM

Tuesday AM, August 29, 1995

The meeting was called to order by chairman Dave Bagby at 1:40 PM. Carolyn Heide secretary.

PICS Proforma

Looking for a volunteer to make one by the November meeting. Still looking ... by about 4:30 Simon Black and Michael Fischer had broken down. 95/202 will be the starter document, it will not be ready prior to letter ballot.

Small Groups worked until 4 PM, results:

Ad hoc power management - Rick White

Based on Rick's paper 95/137R1. Discarded more bit use, receiving and transmitting station must stay awake entire beacon period. ATIMs must be acknowledged. Creation of an IBSS fixes the beacon interval and ATIM window size for the life of the IBSS. ATIM retry MIB parameter needed. IBSS power management is not an option. Simon is going to modify Rick's document and 95/137R2 will be available before the end of the week. Items 1 and 13 above covered by this.

Clause 8/MIB (goal: create updated MIB to match rest of draft) - Bob O'Hara

Rewrote subclause 6.2.6.3 because it introduced new MIB variables and they needed some clean up. Paragraph 3 & 4 have new text about retries for RTS and Data frames. Bob will circulate in document 95/201 this week.

WEP (IV, MSDU vs. MPDU) - Dave Bagby

MSDU vs. MPDU - several strong MSDUs and a couple of fence sitters - result: recommend to change back to MSDU. Dave Bagby will include in letter ballot comments and publish before letter ballot closes (maybe not by much but non-zero), document 95/196. Item 10 above covered by this.

ICV - recommend to change back to CRC-16 from CRC-32. Dave Bagby will include in letter ballot comments and publish before letter ballot closes (maybe not by much but non-zero), document 95/197. Item 17 above covered by this.

*** NOTE:

Late Wednesday afternoon, Michael Fischer read from the minutes of the Jan. 1995 meeting in San Jose, that Kerry Lynn said that CRC 32 is the shortest the government will allow. Dave Bagby will check with Kerry about this.

Non-accept of non-WEP Data frames MIB toggle. Also add to capability field. Michael Fischer to include in his letter ballot comments and circulate via reflector before letter ballot closes, document 95/198.

8. State machine limitations

Document 95/199 will be further updated state machines, maybe barely before the letter ballot.

14. Reassociation and association relationship

Document 95/203 will have a new state diagram for this before the end of the week.

11. Use of the TO_DS bit (con't)

new replacement text: "The TO_DS bit is set to zero for all control and management frames. The TO_DS bit is set to zero for all data frames in an IBSS. The TO_DS bit shall be set to one for all data frames from an associated STA operating in an ESS, with the following exceptions: 1. the Keep_Data_Local MIB parameter is true; or, 2. immediately preceding the data frame the originating STA has sent an RTS and received a CTS from the destination STA which is not an AP."

To be continued ...

3. Random Backoff Time

$cwmin = 4$, $cwmax = 32$. Suggestion that smaller values give better performance, and that that has been decided on before in straw poll (May meeting perhaps).

$cwmin = 31$, $cwmax = 255$. Some suggestion that the smaller the minimum the more the chances that people pick the same one, therefore more collisions.

Selecting 2^{n-1} is a good idea because its easier to implement - OK, initial suggestion changed to $cwmin = 7$, and $cwmax = 63$.

It all comes down to what are the real expected traffic model, how many stations will be contending at any given moment. Existing Ethernet models may provide numbers we can use. In the single station non-contending case smaller numbers maximize throughput. The more contention, the higher the numbers that work.

New sets: $cwmin = 7$, $cwmax = 255$. $cwmin = 15$, $cwmax = 127$.

Simulation results expected tomorrow from Tom Baumgartner.

Meeting adjourned: 5:30 PM

Wednesday, August 30, 1995

The meeting was called to order by chairman Dave Bagby at 8:45 AM. Carolyn Heide secretary.

7. fragmentation vs. dwell boundaries

The purpose of this is the efficiency gain of getting more data into a dwell. There is some feeling that it may not be worth doing, the effort is not worth the gain.

Once you have sent a fragment you can never change its size, so having divided it up to get it into that dwell you are stuck with that size in the next dwell if you need to retransmit. All fragments can be any

size to facilitate this, while fragments (excepting the last fragment of an MSDU) could be of a fixed size without this which is easier for the receiver. However, since fragmentation is based on a MIB variable which can change anyways, everyone needs to build receivers that can receive any size fragment anyways, so maybe this is irrelevant. But this MIB variable cannot change within a single MSDU, so without the dwell/fragmentation issue, the fragments, except the last, can remain fixed size within a single MSDU.

This has the effect of an the MAC due to an intrusion of the PHY into the MAC, i.e. the MAC must know that the FH goes off the air for the hop. This is a religious issue to ISO layer purists. The division of layer function may be more important to this standard due to the one MAC many PHYs situation.

Suggestion: The MAC tell the PHY to send the frame, and the PHY could respond with an error code saying that this won't fit in the dwell. The FH PHY would be the only PHY which returns this error code. The MAC would then not know about dwells at all.

If this was not mandatory, although not an option, it does not burden anyone but the FH PHY implementers for whom there should be some sympathy and some effort to help them use their medium as efficiently as possible.

Straw poll: leave as is - 3; don't do dwell fragment adjustment (fixed value of a_frag_threshold per MSDU - 11. This suggestion will be written up as 95/206.

3. Random backoff time

Excel simulation results, presented by Tom Baumgartner:

	CW					
	7	15	31	63	127	255
STA's contending = 2						
slot	1.9	4.6	10	21	42	86
# collisions	.15	.06	.03	.01	.01	.01
STA's contending = 4						
slot	1	2.6	6	12	26	52
#collisions	.7	.34	.18	.08	.05	.02
STA's contending = 10						
slot	.22	.83	2.2	5	10	21
#collisions	3.5	2.2	1.3	.7	.4	.2
STA's contending = 25						
slots	.02	.2	.7	1.9	4.3	9
# collisions	12	10	7	3.9	2	1

Recommendation: CWmin = 7, CWmax = 255

Straw poll: change D2 text to CWmin = 7, CWmax = 255 (zero based, so 7 means 8 slots) required values - 11; those who don't want to change - 1; abstain - 2. This suggestion will be written up as 95/207.

At <ftp://ftp.cs.toronto.edu/csri-technical-reports/298> is a good technical reference on this subject. Also document 95/183.

9. multirate description inadequate, subclause 6.6

Concerns over last sentence, management frames at any rate? That text is not clear, needs to be fixed.

State transition R41b brings the problem to the forefront - the MAC receiver has been told that the PLCP header was good, but the rate indicated was not a rate that the PHY supported, so the MAC does get the receive data. Therefore no duration field to update the NAV. The MAC knows something is going on and

a guess can be made at what it was, but it doesn't know really. That could be fixed by not allowing multirate, or by a layer violation of putting the first four octets of the MAC header into the PLCP header. There is quite a bit of disagreement with the latter solution.

It is pointed out that the above situation can be created by things other than multirate - the PLCP header can get through while the rest of the frames doesn't.

That problem could be fixed by requiring that if data is going to be sent at non-basic rate must use RTS/CTS and the duration information can be obtained from them. A lot of feeling that that reduces the benefit of using the high rate.

For NAV update during a fragmented transmission, the duration field in the fragments is required. So moving the duration into the PLCP header is the only solution.

While the PLCP header does have length, which indicates how long the medium will be busy for this frame, that is not equal to the duration which covers the entire frame exchange remaining (e.g. the ACK too).

There is a lot of feeling that carrier sense will still work without using the NAV, when it doesn't you get a collision, that's life. There is a suggestion that putting into the standard that the NAV is unreliable if multirate is used. There is agreement that the NAV is not completely reliable all the time anyway, so this is just another case.

Some people feel that the current description of multirate use is just fine. There is disagreement about the action that should be taken in R41b.

There is no agreement on a single solution.

11. use of the TO_DS bit (con't)

Replacement subclause 4.1.2.1.3: "The TO_DS bit shall be set to zero for all control and management frames. The TO_DS bit shall be set to zero for all data frames in an IBSS. The TO_DS bit shall be set to zero in data frames from an associated STA operating in a ESS when the Keep_Data_Local MIB parameter is true. If Keep_Data_Local MIB parameter is false, the TO_DS bit shall be set to one for all broadcast and multicast data frames transmitted by an associated STA. For all other data frames the TO_DS bit may be set to zero or one, at the discretion of the implementation.

A STA, which is not an AP, shall discard all data frames received with the TO_DS bit set."

This subclause may not be the appropriate place for the text on how to use the bit - this subclause is what is the bit. While this may be true, some feel that the purity of what is defined where has been violated many places. This is an editorial concern.

Some people feel that none of this should be added, it is good enough as it is. All implementers need to know is what the bit means, let them decide when to use it.

Without adding a MIB variable, implementation don't have a way to use it.

My point about specification at the receiver.

An opinion is expressed that the Keep_Data_Local parameter may have been not thought out well enough yet - some people think it should be accompanied by a list of stations, others see it as a simple Boolean.

Suggestion that the TO_DS bit be deleted, the address fields can be used.

12. uncompress the TIM bit map

Currently: assumption that the AP assigns lowest possible SID, keeping the map from exploding. TIM contains a bit map of stations indexed by SID. The extra block id indicates only bytes with active bits. This is the compression method. So the receiver must find its bit, which requires more processing for higher SIDs than lower SIDs. Also, standard says if no other bits active, prior backoff not required.

If bitmap not compressed: assumption TIM includes bitmap such that all stations active are in it. So not active, no bits needed map compression can be by cutting off trailing zeros. All stations use their SID as index, same processing regardless of high or low SID. Easy to check if other stations set to determine whether to use prior backoff. AP doesn't have the burden of compression.

If adopted, policy for reusing SIDs is needed, although some say it is needed anyways, especially with the SID collision issue.

People think that there is a compromise, an algorithm that enables compression of large numbers of zero, perhaps leading as well s trailing. A simpler scheme is desired, perhaps all the way to only trailing zero compression is too far.

Straw poll: only trailing compression - 5; leading and trailing - 6; leave it alone - 0

There is a thought that choosing the compression method for the TIM may lead the implementer to choose an SID management scheme which is favorable for this function, so maybe an SID reuse policy is not necessary.

No agreement at this point. Documents 95/208 and 95/209 will describe these two approaches.

15. combining association and authentication.

Suggestion:

"In open system: combine authentication frame 1 and association request frame; and, authorization frame 2 and association response.

Shared key and proprietary systems: combine authentication frame 1 and association request; and, authentication frame 4 and association response."

This removes frames, makes more efficient, changes no functionality. Some feel that for that small an overhead removal it is not worth the trouble to change now. Others that it is a significant overhead, especially in an open system. Yes, but is that the case for which we optimize?

We know that each authentication and association have to be done independently (for IBSS operation authentication only, and for pre-authentication, authentication only), and have to sometimes be both done. So this proposal means that we need to leave combined and independent methods. Having two ways of doing the same thing is maybe not a good idea.

Straw poll: change to this - 2; leave as is - 10.

So, there we are.

16. what to do if station leaves, gets disassociated, then returns and it thinks it still owns the SID.

If STA 1 receives a class 3 Frame from STA 2 which is not associated with STA 1, STA 1 shall send a Disassociation frame to STA 2.

Straw poll to adopt: 7,0,0. Document 95/210 will describe this.

Relevant to this, is the changing of the reassociation frame to be a class 3 frame which happened earlier this week (Item 11, document 95/203).

18. TSF timer algorithm is screwed up.

Used to have a distributed time synchronization method - receive a time, if different from yours take the difference. This converges quickly, regardless of which station is sending beacons. Now we have - adopt the highest time seen, so if you have a high time and don't send it often, you will take a long time to synchronize.

There is a sentiment that the original scheme of coalescing by taking the difference worked very well, and has worked in other protocols so we should go back to it.

This change was made trying to avoid having time moving backwards for a station. However, time can move backwards when a station joins an infrastructure network - you adopt the time given to you when you associate.

There is agreement that 8.1.4 may be OK, but subclause 8.1.1.2 is what is screwed up. There is not agreement on what to do to fix it.

The rule may be: A station joining an IBSS for the first time shall not transmit a beacon until it hears a beacon or probe response. [straw poll: 10, 1, 0] After that it shall transmit beacons with its TSF timer.

When and what to set the time stamp to in a beacon? A sentiment is expressed that it is not adequately specified, but that the MIB variables are there to specify it properly. Further examination of subclause 11.3.3.1.2 reveal that the timer reference in the MAC doesn't work - either the PHY or MAC needs fixing. PHY service definitions are inadequately and inconsistently specified.

No resolution achieved. More research needed.

19. how to facilitate key change

95/187 deals with this. Hooks for things that are outside of the standard - some way is needed to identify the key which is in use so that things don't get lost when the key changes (things that were outstanding, for example retransmissions that were already encrypted with the old key, or stations that were asleep when the key changed). So a key id allows you to say what key this message was transmitted with. This could be put in the IV pad byte, suggestion is using the two lower bits as the key id. This is for the shared key case which is trying to synchronize over a wide range of stations, rather than a bilateral key mapping case.

Document 95/211 will describe this.

20. should simultaneous MSDU limit be different for AP vs. non-AP, and what happens if maximum exceeded?

Suggestion that for non-AP stations it should be reduced to one or maybe two. Currently subclause 6.4 says six for all stations.

Should we make stations which are AP capable and not - we probably want to make just stations, period. You take your station and put it into a machine that is or is not an AP.

Straw poll: Reducing below six - 1; leaving as is - 12.

What happens if six exceeded? Straw poll: do we need to specify this - 8, 3.

There follows a discussion of what do when there is room for MPDU but not all of the MSDU, and whether or not to ACK. Repeated Straw poll: do we need to specify this - 4, 6.

Needs to left for individual letter ballot comment.

21. beacon response to poll vs. power save mode

Station polls AP and is waiting to be sent data. New beacon arrives with bit set again. Does the station poll again? Subclauses 8.2.1.5 and 8.2.1.7 do cover this. It also clearly states that once you have sent a poll and the AP has acked it, you must stay awake until you get something from the AP or you get a TIM that does not have your bit set. So this is not a problem.

The only interesting case is how fragmentation and transmission of ATIMs interact in an IBSS. If you receive some fragments of an MSDU then a beacon, do you have to get an ATIM again to tell you to stay awake? Most people say yes, others say stay awake until either you get them all or the data lifetime expires.

Further input is left to individual letter ballot comments.

4. & 6. MIB out of date and Connection Stuff

These seem to be decision that were made, but the changes didn't get into the draft.

Decisions made that didn't get into D2: connection stuff Document 95/212 will describe this; WEP MSDU vs. MPDU covered in items above; format and figures in subclause 4; TIM diagram out of date (4.3.2.1); Figure 6-10 correct diagram which is in 95/174.

Small groups for status codes, MIB stuff and other things meet at 8:30 AM tomorrow, for now there are a lot of papers to create overnight, so adjourn.

Meeting adjourned: 4 PM

Thursday AM, August 31, 1995

The meeting was called to order by chairman Dave Bagby for a few minutes to decide that what should be done for a couple of hours until the full working group meeting was work on the little papers that people agree they would produce. So they did that.

Meeting adjourned: around 9 AM