Collected comments on Section 5 of draft standard D1 (PART 2)

5.2.6.1. & 5.2.6.2.	Fischerma:Basic Access & Backoff Procedure	T	Backoff mechanism must be changed to state that "stations in backoff should count backoff time whenever the medium is sensed free," as opposed to only after DIFS. Alternative solution: transmitters must backoff on initial attempts.	Network capture Is more likely with the D1 proposal than it ever has been with 802.3 because of the procedure described in D1 sections 5.2.6.1. and 5.2.6.2. For example, if four stations are involved in two higher-level conversations, then the first winner of a contention period will then keep the medium for as long as he has traffic to transmit. This is because he will never find the medium busy (since whenever he has completed a transmission, the medium should be free again) unless he must by default backoff for each new transmission - but I do not find such wording anywhere in the document. (See section 10.3.3.2. Carrier Sense/Clear Channel Assessment Procedure) The winner of the first backoff will transmit an uninterrupted stream of traffic, since the loser is not allowed to count down his backoff until the winner has no more traffic to transmit. This is because the loser is not allowed to count down his backoff unless the medium has passed DIFS with no traffic, but at the end of each DIFS, just when the loser would start his backoff counter, there will be a new frame from the winner. Effectively, the winner will have captured the network - the loser of the contention will have chosen a non-zero backoff value, and he may only count down when the medium is NOT busy following a DIFS. But the medium will always be busy following DIFS as long as the winner has traffic to transmit! At least in the 802.3 case, the loser was allowed to count down his backoff even if the network was busy. He then had a chance, after some later IFS time, to attempt to come back in and win the contention back from the original winner. I vaguely remember in a proposal that all intial TX attempts must use an initial backoff, but I do not see that anywhere in the D1 document.
5.2.6.2	Bob O'Hara	E	Replace "selecting" with "computing."	
5.2.6.2	Bob O'Hara	E	replace "A station that has just transmitted a frame" with "A station that has just completed transmission of a frame", delete the comma and move "to the medium" after "access" in the paragraph after the figure.	Better usage, clarity.
5.2.6.2	David Bagby	E	A station that has just transmitted a frame and has another frame ready to transmit (queued), shall perform the backoff procedure. This requirement is intended to produce a level of fairness of access amongst STAs to the medium.	See imbeded comments and annotations
5.2.6.2	A. Bolea	T		Clearly state that the backoff timer is only decremented after the medium is idle for a slot time and not continuously.
5.2.6.2	Bob O'Hara	T	add "and placing that value into the Backoff Timer" to the end of the first sentence of the second paragraph.	A method for initializing the Backoff Timer must be described.

5.2.6.2	Geiger	T	~~The advantage of this approach is that stations that lost contention will defer again until the next DIFS period, and will then likely have a shorter backoff delay than new stations entering the backoff procedure for the first time.~~	This statement is not true unless the Random() function uses some inverse weighting of the CW value. Stations entering the contention period for the first time will have a better chance of winning access than the units which have already backed off because their CW is smaller than backed off units and they have a higher probability of picking a smaller number than the stations that have already contented once. What this algorithm really does is provide a means for reducing collisions in a congested state where lots of stations are trying to access the medium. This algorithm helps reduce the number of units picking the same slot time. It does not tend toward fair access on a first come, first serve basis.
5.2.6.2	Greg Ennis	T	remove final paragraph	explanations are not necessary within the standard text
5.2.6.2	P. Brenner	T	It should be specified that for the purpose of Backoff Procedure, the Contention Free Period is to be considered as "busy" medium, i.e. Backoff Timer does not decrement, even when the medium is sensed free.	This will reduce the probability of collisions immediately after the Contention Free Period.
5.2.6.2	Renfro	T		State that backoff timer is decremented in steps of slot time. Need to ensure that stations which lose contention during random backoff will begin transmission on integer slot time next time around.
5.2.6.2	Rick White	T	Assuming that the backoff timer is integer multiples of the slot times, the backoff timer should be decrement after each slot time in the contention window when the medium is not busy. When the backoff timer reaches zero, the STA should access the medium.	
5.2.6.2	Rick White	T	The statement: "A station that has just transmitted a frame and has another frame ready to transmit (queued), shall perform the backoff procedure." is not true.	A station does not have to perform the backoff procedure when it has received an ACK for a fragment of a fragmented MSDU and has an additional fragment for the same MSDU to send. This must be corrected.
5.2.6.3	A. Bolea	E		**"CW will be greater than one.."** should be reworded to reflect correct CW as defined insection 5.2.5.
5.2.6.3	Greg Smith	E	ACK_RE-TRANSMIT_counter and ACK_RE-TRANSMIT_Limit should be : DATA_RE-TRANSMIT_counter and Limit	If this is not an editorial error, then much more explanation is requir
5.2.6.3	McKown	E	If after ... > If, after ...	phrase needs commas at both ends
5.2.6.3	Rick White	E	This section should come after Section 5.2.6.4 which describes the use of RTS/CTS.	
5.2.6.3	Wim Diepstraten	E	Add a paragraph as follows: Stations that receive an RTS frame, but do not sense a busy medium (Data frame) after a CTS-Timeout period can reset their NAV to the previous value. Last paragraph, middle sentence: Change "CW will be greater then one" into: "CW will be greater then Cwmin"	The last paragraph does again suggest that Cwmin is one rather then a much bigger value (for instance 16 or 32).
5.2.6.3	bdobyns	T	RTS_RE-TRANSMIT_LIMIT and ACK_RE-TRANSMIT_LIMIT are not defined elsewhere. Either put them in the MAC MIB or use MAC MIB parameters.	maybe these should be aRTS_Retry_Max and aDATA_Retry_Max?
5.2.6.3	Bob O'Hara	T	Replace the two retransmit limits with a single limit.	It is not clear why the MAC should try harder to deliver a frame in one case than another. No mechanism is described to initialize the two limit counters and how to handle interaction between them. Two limits are unnecessarily complex.

5.2.6.3	David Bagby	T	Many circumstances may cause an error to occur in a RTS/CTS exchange. For instance, CTS may not be returned after the RTS transmission. This can happen due to a collision with another RTS or a DATA frame, or due to interference during the RTS or CTS frame. It can however also be that CTS fails to be returned because the remote station has an active carrier sense condition, indicating a busy medium time period. this section appears to crate a different retry limit for RTS than non-RTS cases. I don't agree with nor see the usefulness of this. RTS frames should be retried the same as other frames. If after an RTS is transmitted, the CTS fails in any manner within a predetermined CTS_Timeout (T1), then a new RTS shall be generated while following the basic access rules for backoff. Since this pending transmission is a retransmission attempt, the CW shall be doubled as per the backoff rules. ~~This process shall continue until the RTS_RE-TRANSMIT_Counter reaches an RTS_RE-TRANSMIT_Limit.~~ ~~The same backoff mechanism shall be used when no ACK frame is received within a predetermined ACK_Window (T3) after a directed DATA frame has been transmitted. Since this pending transmission is a retransmission attempt the CW will be greater than one as per the backoff rules. This process shall continue until the ACK_RE-TRANSMIT_Counter reaches an ACK_RE-TRANSMIT_Limit.~~	See imbeded comments and annotations
5.2.6.3	David Bagby continuation	T		
5.2.6.3	Geiger	T	RTS_RE-TRANSMIT_Limit RTS_RE-TRANSMIT_Counter ACK_Window (T3) CTS_RE-TRANSMIT_Limit CTS_RE-TRANSMIT_Counter CTS_Timeout(T1)	I assume that these values need to be defined somewhere, maybe the MIB, can't find them there.
5.2.6.3	Greg Ennis	T	remove first two paragraphs	explanations are not necessary within the standard text
5.2.6.3	Rick White	T	Must resolve editor's comments relating to CTS_Timeout (T1), RTS_RE-TRANSMIT_Limit, ACK_RE-TRANSMIT_Limit and ACK_Window (T3) and any requirement for interaction between RTS and ACK retransmission.	
5.2.6.3	Okada	T Approve	No T1 and T3 are defined	They are defined by each PHY

5.2.6.3, 4	Geiger	T	Okay, now I know how RTS and CTS are exchanged, when do you send the data frame which generated the RTS/CTS exchange, during the SIFS slot or the first DIFS slot or what? I suspect that a station should send the data frame involved in a RTS/CTS exchange in the SIFS slot as well as the ACK, CTS, and Data Fragments.	Guys, we have got to think of all these things and then document them. - I found it... its in figure 5-9. Good job hiding this one. I just finished playing Myst and this clue was tougher than any Myst clue to find. Section 5.2.4.1 should be re-written to include RTS/CTS data frame and data fragments!
5.2.6.4	C. Heide	e	2nd paragraph, second sentence, strike the first word "a".	
5.2.6.4	C. Heide	e	figure 5-9, explain T1 and T3.	
5.2.6.4	Tim Phipps	E	*Remove:* "In the absence of a PCF ... reset the NAV".	This is not true. Data and ACK frames also carry duration information and update the NAV.
5.2.6.4	A. Bolea	T		In second paragraph, need to clarify that destination station of RTS does not update its NAV.
5.2.6.4	A. Bolea	T		Accuracy of NAV should be in units of milliseconds.
5.2.6.4	Bob O'Hara	T	Replace "X ns" with "one bit time."	Simpifies timing requirements. (See comment on 4.1.2.2)
5.2.6.4	C. Thomas Baumgartner	t	state NAV internal state accuracy as - 0 + 1 microsecond.	A longer NAV will not cause protocol errors but a shorter NAV counter will.
5.2.6.4	Dean Kawaguchi	T	Setting the NAV Through Use of RTS/CTS Frames (3rd paragraph) Maintenance of the NAV shall consist of an internal state accurate to ~~X nS~~ 1 microsecond of the busy/free condition of the medium...	Standard should not have a TBD. The uncertainty of other timing factors such as propagation delay is on the order of 1 microsec. The allowed error in the NAV should not be any more stringent.
5.2.6.4	Fischer, Mike.	T	The X nanoseconds in the 3rd paragraph needs to be quantified. My recommendation is a value of 1000ns (1 microsecond, the same resolution as the TSF timer).	Many things in this MAC are done to microsecond resolution, so there is no simplification to using a NAV resolution coarser than 1 microsecond. Given the response times in the existing PHY specifications, there appears to be no benefit to a finer NAV resolution than 1 microsecond.
5.2.6.4	Geiger	T	Duration Field value determination is not defined. See section 4 for RTS and CTS frame structure says to see section 5.xx for duration field explanation. (I believe programmers call this an infinite loop).	Don't worry about this, just throw RTS/CTS out!
5.2.6.4	Joe Kubler	T	X nS should be defined.	
5.2.6.4	Joe Kubler	T	strike "are the only events that"	NAV should be set to protect ack on directed data MPDU
5.2.6.4	Joe Kubler	T	dwell should be "dwell or superframe"	in presence of PCF AP, should fragment based on superframe time (not just FH hop dwell time).
5.2.6.4	John Hayes	T	TBD	Accuracy of X ns needs to be defined.
5.2.6.4	Mahany	T	Required Accuracy of NAV timer must be inserted. 1usec +/- 25 ppm is appropriate.	Required for Interoperability
5.2.6.4	McKown	T	para 3: NAV precision specified as "X nS"	typo
5.2.6.4	Miceli	T	NAV accuracy needs definition	not defined
5.2.6.4	Paul Pirillo	T	Insert: The NAV has a value in the range {XX..YY} that is an integer multiple of the slot time	I am unclear as to the upper limit for NAV and whether NAV is a multiple of some other time period (such as slot time) or whether NAV can take on any value in the valid range. The text I suggest is just an example of how to resolve my concerns. I will accept any text that defines these properties of NAV.
5.2.6.4	Paul Pirillo	T	Insert: The NAV has a value in the range {XX..YY} that is an integer multiple of the slot time	I am unclear as to the upper limit for NAV and whether NAV is a multiple of some other time period (such as slot time) or whether NAV can take on any value in the valid range. The text I suggest is just an example of how to resolve my concerns. I will accept any text that defines these properties of NAV.

5.2.6.4	Renfro	T		Accuracy of NAV counter of X ns must be better defined. Nanosecond timing is not necessary in this network. If allowable inaccuracy grows to several usec, must be included in slot time since it will result in error in starting DIFS timer.
5.2.6.4	Rick White	T	The reception of Data and ACK frames can also set the NAV to a non-zero duration.	
5.2.6.4	Rick White	T	Data and ACK frames also contain a duration field.	
5.2.6.4	Rick White	T	Must define the internal state accuracy for the NAV.	Not defined.
5.2.6.4	Stuart Kerry	T	Setting the NAV Through Use of RTS/CTS Frames (3rd paragraph) Maintenance of the NAV shall consist of an internal state accurate to ~~X nS~~ 1 microsecond of the busy/free condition of the medium...	Standard should not have a TBD
5.2.6.4	Tom T.	T	Change 'X nS' value to 1 µsec. Remove word 'a' from third line second paragraph 'All STA receiving a these...'	Need a real number here. 1 µsec seems reasonable considering the size of the DIFS.
5.2.6.4	Wim Diepstraten	T	Stations should set the NAV to the received "Duration" field only when the "Duration" value is greater then the current NAV value.	Stations could have already received other RTS and CTS information (from a neighbouring BSS) that has already set the NAV to a larger value then the new "Duration" value.
5.2.6.4	Wim Diepstraten	T	The "Duration" field in Data and Ack frames should also be interpreted by all stations, and they should update their NAV accordingly. Also section 5.2.10 should be updated to reflect this procedure.	The text in these sections does not reflect the changes that occured by the fragmentation. The "Duration" field in the Data frame should be specified similarly as the RTS function, while the Duration field in the Ack should be specified similarly as the CTS function.
5.2.6.4	Greg Smith	T/E	X nS	The value of X is fundamental to the operation of the system
5.2.6.4.	Fischerma:Seting NAV through use of RTS/CTS frames	T	In the absence of a PCF, reception of RTS, CTS, DATA and ACK frames may all set the NAV to a non-zero duration in certain circumstances.	Note that D1 text fails to include DATA and ACK frames that are part of a fragmented MSDU exchange as being capable of setting NAV to non-zero value.
5.2.6.4.	Fischerma:Setting NAV through use of RTS/CTS frames	T	Maintenance of the NAV shall consist of an internal state accurate to 16 microseconds of the busy/free condition of the medium.	Xns resolution of NAV is not specified. 16 microseconds would satisfy the update rate of the CCA information delivered by PHYs.
5.2.6.4.	P. Brenner	T	The Duration field of the CTS should be copied from the RTS (without any need for further calculation), so its definition should be: On the RTS it is the time from the end of the corresponding CTS to the end of the ACK, and on the CTS it is the time from the end of this message to the end of the ACK.	The CTS (and ACK) response is the more time-critical portion of the whole MAC implementation, so the amount of calculations in this portion should be reduced. Outsider stations (calculating the NAV) are idle, so the calculations overhead should be there.
5.2.6.4.	P. Brenner	T	The accuracy of the NAV timer should be on the range of 1 microsecond	The PHY specifications are in microseconds, so there is no point of having a NAV more accurate than that.
5.2.6.5	Bob O'Hara	E	Replace "IFS" with "SIFS"	typo?
5.2.6.5	Bob O'Hara	E	Replace "will" with "shall" in second paragraph	Proper standard language
5.2.6.5	Bob O'Hara	E	Replace "IFS" with "SIFS" in caption to figure 5-10	
5.2.6.5	Bob O'Hara	E	Replace "will" with "shall" in tenth paragraph	Proper standard language
5.2.6.5	C. Heide	e	second paragraph, first sentence remove the word "either".	
5.2.6.5	C. Thomas Baumgartner	e	in first sentence change from (IFS) to (SIFS) in 3rd paragraph change from IFS to SIFS Change in title of Figure 5-10 from IFS to SIFS	typo

5.2.6.5	Geiger	E	The Short Inter-Frame Space (SIFS) (beginning of the section the IFS abbreviation is wrong)	Clarity
5.2.6.5	Iwen Yao	E Approve		Dwell Time is used in this section but it is not defined. Please define.
5.2.6.5	Mahany	E	Define Dwell Time Prior to this discussion.	Readibility. Concept of a dwell time has not been introduced at this point. Superframe boundary may provide same limitation. as dwell time (5.3.1) See 5.5 for dwell time definition.
5.2.6.5	Geiger	T	When a station has transmitted a frame other than a fragment, it does not have priority to transmit on the channel following the ACK for that frame	What the hell does not have priority mean? Can't use the SIFS frame or must use the backoff algorithm, or start with a CW of 2 or what? Why not say the station must wait for the normal contention period before it can again access the channel.
5.2.6.5	Greg Ennis	T	replace "until ... MSDU" with "until it has sent all fragments of an MSDU and received their corresponding ACKS, or until it failed to receive an ACK for a specific fragment".	Station must contend if it fails to receive an ACK for a fragment
5.2.6.5	McKown	T	para 3: guidlines > rules	not optional
5.2.6.5	McKown	T	para 13 et seq.: a limit on the number of fragment retransmission attempts in the absence of acknowledgement, without the use of RTS & CTS, should be established (analogous to RTS_Retransmit_Limit, which applies with RTS & CTS).	oversight?
5.2.6.5	Renfro	T		In paragraph 7, if source station receives ack but does not have time to transmit next fragment and receive ack before hop, it must not only contend for channel after hop settling time, it must use random backoff procedure.
5.2.6.5	Renfro	T		Last paragraph. Frames not requiring ack (e.g., broadcast/multicast from AP) should not be fragmented. Probability of success will be higher if they are transmitted in entirety since no ack to indicate failure.
5.2.6.5	Rick White	T	¶ 8, Fragment retransmission: Change "If the source station does not receive an acknowledgment frame, it will attempt to retransmit the fragment at a later time (according to the backoff algorithm). When the time arrives to retransmit the fragment, the source station will contend for access in the contention window." to "If the source station does not receive an acknowledgment frame, it will attempt to retransmit the fragment according to the backoff procedure. The CW shall increase exponentially after every retransmission attempt for any fragment for a given MSDU, up to a maximum value CWmax.	
5.2.6.5	Tom T.	T	Add to the last paragraph: The spacing between fragments of a broadcast/multicast frame shall be equal to the SIFS period.	The case of a broadcast fragment burst is unique and must be fully specified. From the implementation point of view it would be easier to make this a PIFS time, however this should be so only if it is not possible for a STA to send a fragmented broadcast during the contention free period of a superframe.
5.2.6.5	Okada	T Approve	If the source station does not receive an acknowledgement frame, it will attempt to retransmit the fragment at a later time (according to the back-off algorithm). How long does the source station have to wait, T3 or SIFS?	T3
5.2.6.5.	P. Brenner	T	It should be specified whether contiguous Fragments of MSDUs that do not require acknowledgment are sent with SIFS between them or not.	Is not clear from the draft.
5.2.6.6	A. Bolea	E		In Figures 5-11& 5-12, the "RTS" needs to be blocked off.
5.2.6.6	Bob O'Hara	E	Add box around opening RTS in figure 5-11	
5.2.6.6	Bob O'Hara	E	replace "Frame" with "Fragment" in the last sentence before figure 5-12	
5.2.6.6	Bob O'Hara	E	Add box around opening RTS in figure 5-12	
5.2.6.6	C. Heide	e	figure 5-11 and 5-12 are missing boxes around RTS frames	

5.2.6.6	C. Thomas Baumgartner	e	Put x through or shade differently the NAV(ACK 1) in Figure 5-12. Figure 5-12 missing a box around RTS	since the discussion says there was no ACK 1 to create that NAV update. Typo.
5.2.6.6	Jim Panian	E	Specify that each fragment is transmitted after waiting SIFS.	The draft states that "the source station will transmit all fragments of the MSDU without releasing the channel as long as there is enough time left in the dwell time". Does this mean that there is no SIFS between fragments?
5.2.6.6	Jim Panian	E	Change the last sentence of the second paragraph to read "Each fragment and ACK acts as a virtual RTS and CTS for the next fragment to come."	The text is ambiguous regarding the applicability of the duration field for fragments and ACKs.
5.2.6.6	Jim Panian	E	Remove the NAV (ACK 1) from "Other" from the figure "RTS/CTS with Transmitter Priority w/ missed ACK."	The figure is incorrect in showing the NAV being set by ACK 1 when ACK 1 is never sent.
5.2.6.6	Jim Panian	E	Place RTS in the two figures.	RTS is not within a "box" of the following two figures: • RTS/CTS with Fragmented MSDU • RTS/CTS with Transmitter Priority
5.2.6.6	Rick White	E	This section should be moved to be after Section 5.2.7	The basic RTS/CTS function should be introduced before addressing RTS/CTS for fragmentation.
5.2.6.6	Tom T.	E	Figure 5-11 errors: Correct the NAV bar for fragment 1 to start from the end of Fragment 1, not from the end of ACK1. Same for NAV bar of Fragment 2.	
5.2.6.6	Okada	E Approve	Clarification. Does NAV have a count-down timer which is defined by PHYs?	
5.2.6.6	A. Bolea	T		The way the duration field is defined in fragments, a station will need to hold on to the duartion value until the ACK is complete before using it as the NAV. I recommend that we redefine the duration in each fragment so that it corresponds to the time from the end of that fragment to the ACK for the following fragment. This makes the RTS and Fragment NAV processing identical. That is, any time a valid message(RTS,CTS, ACK or Fragment) is received, a station sets the NAV to the duration value. CTS and ACK processing is also identical because the duration field is calculated as the duration field from the proceeding RTS or Fragment minus a fixed offset. Figure 5-11 needs to be updated to reflect this.
5.2.6.6	A. Bolea	T		In last two sentences, it is stated that a station which has not received an ACK should wait until NAV has expired before attempting re-transmission. I recommend that the station be allowed to re-transmit after a DIFS plus random backoff as it would do normally for any re-transmissions. Figure 5-12 needs to be updated to reflect this.
5.2.6.6	Joe Kubler	T	strike sentence "The source station must wait until the NAV (Fragment 1) expires before attempting...	While this adds a little to fairness of access, it wastes a potentially large amount of bandwidth
5.2.6.6	Renfro	T		Define duration field to be the time from end of current frame till end of next anticipated ack in all cases. This makes processing consistent whether the duration information is in a RTS, CTS, Data or ACK frame.

5.2.6.6	Renfro	T		Transmitting stations should not maintain NAV. When ACK is not received, transmitting station should try to reaccess the channel beginning after anticipated ACK would have been received. This is the same time that stations getting NAV information from the CTS will begin to access the channel.
5.2.6.6	Rick White	T	Must define how a STA makes a decision to use RTS/CTS for a fragmented MSDU. I assume that if the fragment size is greater than RTS_Threshold, RTS/CTS is used.	Not defined.
5.2.6.6	Rick White	T	The duration field in the Data and ACK frames shall be used to update the NAV even if the transmission did not start with an RTS/CTS exchange.	This is especially useful for a multi-fragement MSDU that does not use RTS/CTS.
5.2.6.6	Rick White	T	If RTS/CTS is used for the initial transmission of a fragmented MSDU, RTS/CTS will be used for retransmission of any fragments of the MSDU.	
5.2.6.6	Tim Phipps	T	Figure 5-11. 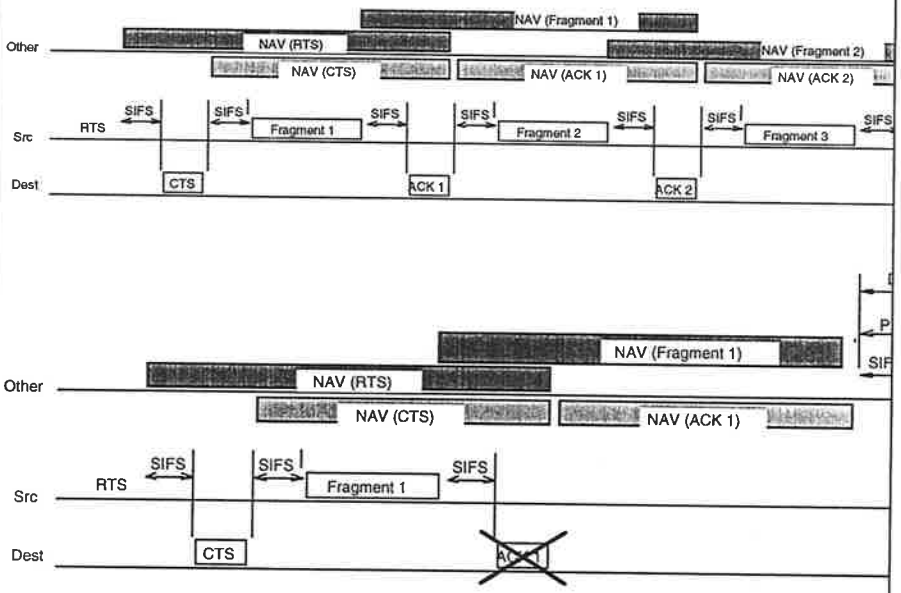	The NAV should always be updated at the end of a received packet.
5.2.6.6	Wim Diepstraten	T	The figures and description in this section should be updated to reflect the general definition of the "Duration" field when it is used in a Data and Ack frame. The Duration field in a Data frame should specify the time from the end of the data frame until the end of the Ack of the subsequent fragment (RTS function). The Duration field in a Ack frame should specify the time from the end of the Ack frame until the end of the subsequent Ack frame (CTS function).	This definition is similar to the Duration definition for an RTS and CTS frame. The same CTS_Timeout mechanism could be used to reset the NAV when a subsequent fragment is not immediately send as result of a Ack failure.

5.2.6.6 and 5.5	Iwen Yao	T Approve		It is not clear that whether RTS/CTS is required if the same MSDU has to recontent for the medium for any reason. e.g. If the dwell time expired before all the fragments are sent. Please clarify. It seems reasonable to explicitly require the use of RTS/CTS in this situation if it is used to set up the transmission of the MSDU.
5.2.6.6.	Mahany	E	For improved clarity in second paragraph, may wish to insert sentence: "Fragment 2 and ACK 2 will update NAV to indicate busy until end of ACK 3. ", prior to last two sentences.	Readabilty
5.2.6.6.	Fischerma:RTS/CTS usage with Fragmentation	T	First paragraph of section: The following is a description of using RTS/CTS for a fragmented MSDU. The RTS/CTS frames define the duration of the first frame and acknowledgment. The duration field in the data frame specifies the total duration of the subsequent ACK frame, the next frament and the next ACK frame and the duration field in the acknowledgment frames specifies the total duration of the next fragment and acknowledgment. This is illustrated in Figure 5-11.	NAV should be updated at end of frame received to avoid hidden node problem. Current description tries to avoid hidden node problem by relying on storing duration field from DATA frame until NAV timeout, and then reloading NAV at that point. This method is unacceptable, since it is inconsistent with RTS/CTS NAV update scheme.
5.2.6.6.	Fischerma:RTS/CTS usage with Fragmentation	T	Text and diagrams should be updated to convey the following directive: NAV counter shall be updated with new duration field information at the end of the successfully received frame from which the duration field was parsed. [This implies that duration field information for DATA and ACK frames of MSDU fragments must be different - DATA duration must include this ACK, next DATA, next ACK, ACK duration field should include next DATA, next ACK.]	It is unclear from the text and the diagrams when the NAV should be updated. Should the update for DATA frame duration field inforamtion occur at the end of the DATA frame, or at the end of he ACK frame, or at the end of the curent NAV count, assuming that the ACK frame is not received first? I vote for: update NAV at end of frame that contains duration and is successfully received, since this is consistent with current description of NAV updates for RTS and CTS frames.
5.2.6.6.	Fischerma:RTS/CTS usage with fragmentation	T	Last paragraph of section: The source station must wait until the ACK timeout before attempting to contend for the channel after not receiving the acknowledgement.	Note that D1 wording implies that source station maintains a NAV according to its own transmissions! NAV update policy elsewhere in D1 makes no mention of NAV updates in response to own transmissions. Also, it is not clear that even if NAV was updated during say, ACK frame of fragmented MSDU exchange at the DATA frame sender, that the DATA frame sender would somehow be allowed to ignore the NAV information in order to send the next DATA fragment. Therefore, wording should reflect accepted transmitter behavior by obeying ACK timeout in order to determine when to begin contending for channel again.
5.2.7	A. Bolea	E	Last sentence "frame and **an** SIFS gap period. No regard shall be **give** " should be "frame and **a** SIFS gap period. No regard shall be **given**"	
5.2.7	Bob O'Hara	E	Change last sentence of third paragraph to: "The value zero shall be used to indicate that all MPDUs shall be delivered with the use of RTS and CTS."	Clarity
5.2.7	Bob O'Hara	E	delete "gap" and replace "give" with "given" in last paragraph	
5.2.7	C. Heide	e	third paragraph refers to "LME" which is undefined	
5.2.7	Joe Kubler	E	last sentence "shall be give to" should read "shall be given to "	
5.2.7	Renfro	E		Add 'LME' to list of acronyms.
5.2.7	Bob O'Hara	T	Clarify or delete paragraph four.	It is ambiguous
5.2.7	C. Heide	t	remove fourth paragraph	a STA's RTS_Threshold has no control over incoming frames - the sending STA's RTS_Threshold controls whether it uses RTS/CTS or not, and is it does the receiving STA must adhere to that regardless of its own RTS_Threshold. Therefore this parameter does not control direction.

5.2.7	C. Heide	t	last paragraph - carrier sensing should be done before any access to the medium.	hidden stations, stations with varying coverage distances and unsymetrical rx/tx distances will cause many instances of STAs accessing the medium when they shouldn't. Collisions can be minimized by carrier sense before any transmit.
5.2.7	C. Thomas Baumgartner	t	change 1st sentence of 2nd paragraph to "STA shall use an RTS/CTS exchange for directed frames according to the state of attribute (NEED NAME OF THIS ATTRIBUTE) with values of never and when the length of MPDU is greater than the length threshold indicated by the RTS_Threshold attribute."	5.2 says use of RTS/CTS can be set to always, never, or when MPDU greater than threshold. RTS_Threshold value can't take care of never state so an RTS/CTS attribute is required. Personally, I'm not sure that never should be allowed because of the implications for operation in overlapping BSA's
5.2.7	C. Thomas Baumgartner	t	Need discussion of affect on overlapping BSA in same channel of not sensing medium before CTS	MUST have simulation of this affect to know if this is good design.
5.2.7, 3rd paragraph	Fischer, Mike.	T	Add sentence at end ÔThe value 2304 shall be used to indicate that no MPDU shall be delivered with the use of RTS/CTDS.Ó	completeness
5.2.7.1	Renfro	E		Combine with 5.2.7. Inappropriate to have only a single subheading. Add reference to figure 5-13 in text.
5.2.7.1	David Bagby	T	**Figure 5-13: Directed Data/ACK MPDU**	See imbeded comments and annotations
5.2.7.1	Joe Kubler	T	data should set duration to protect the ack	in a busy network, the number of missed acks could get quite large without this. it really adds no cost to bandwidth since (as fig 5-13 shows) other stations should defer until after a DIFS following the ack. This would still allow the use of short directed frames even in BSAs that are using RTS/CTS in an efficient manner
5.2.7.1	Rick White	T	Figure 5-13 should be modified to show that the data frame is transmitted at some point during the contention window, not after DIFS.	STA must select a window in the contention window after DIFS.
5.2.8	Bob O'Hara	E	replace "STA's" with "STAs"	Proper usage.
5.2.8	Bob O'Hara	E	replace "on" with "for" in the last paragraph	Proper usage.
5.2.8	Tom T.	E/T	How does a STA decide whether to send a broadcast STA to STA or through the AP? What does the AP do with a broadcast frame it hears from a STA to STA transmission?	
5.2.8	A. Bolea	T		Broadcast/multicast messages should not be fragmented since we don't all receiving stations trying to ACK the fragments. In Infrastructure networks, the broadcast message from the STA to AP should go up the AP as a directed message(it could be fragmented!). The AP would then transmit this entire broacast message without fragmentation.
5.2.8	bdobyns	T	Change "There is no MAC level recovery ... except for those frames sent via an AP." to "... those frames sent *to* an AP"	frames from an AP cannot be recovered if lost.
5.2.8	Bob O'Hara	T	Clarify paragraph three.	It is ambiguous

5.2.8	C. Heide	t	broadcasts and multicasts coming from a STA should be sent without RTS/CTS and ACK. It should be the responsibility of the AP to retransmit them onto the DS. OR STA should send all broadcasts to the AP only if there is one. The AP must then retransmit them within the BSS and onto the DS if the BSS is part of an ESS.	this section proposes that a STA must transmit all broadcasts and multicasts twice - once to the STAs in its BSS and once to the AP so that the AP can distribute them throughout the ESS. This is an unreasonable request of a STA. A STA should not have to know if there is an ESS, or if there is a portal somewhere through which its broadcast must be sent to wired STAs - it should just transmit a broadcast when it needs to do so.
5.2.8	John Hayes	T	TBD	Broadcast and Multicast frames may be fragmented.
5.2.8	Renfro	T		Second paragraph is only true if To DS bit is set. While it is probably a good idea, we have done nothing to preclude individual stations from sending broadcast/multicast messages to everyone. We have also not precluded STA to STA communications without using the AP in an infrastructure network.
5.2.9	A. Bolea	E	ToAP needs to be changed to ToDS	
5.2.9	Fischer, Mike.	E	Change ÒToAPÓ to ÒToDSÓ	correctness
5.2.9	Geiger	E	Can't find Ack_timeout in MIB table	Helps to define it
5.2.9	Greg Smith	E	references 'ToAP' bit should be 'To DS' bit	
5.2.9	Renfro	E	Change 'To AP' to 'To DS'	
5.2.9	Tim Phipps	E	*Replace: "*ToAP" *with* "ToDS".	This has equivalent functionality for the purpose of this section. The "ToAP" bit has been removed.
5.2.9	Bob O'Hara	T	add "without receiving an ACK frame" after "time" in the second paragraph	It is unclear what the purpose of the timeout is.
5.2.9	C. Heide	t	the medium should be sensed before ACKs are transmitted.	Not sensing the medium could cause a collision which destroys the ACK, causing a retransmission which would have resulted anyway had the medium be sensed - no difference. However not sensing the medium causes the other transmission to be corrupted also, which would not have happened.
5.2.9	C. Thomas Baumgartner	t	Need discussion of affect on overlapping BSA in same channel of not sensing medium before ACK	MUST have simulation of this affect to know if this is good design.
5.2.9	David Bagby	T	The Source STA shall wait an Ack_timeout amount of time before concluding that the MPDU failed. This policy induces some probability that a pending frame in a neighboring BSA (using the same channel)	See imbeded comments and annotations
5.2.9	Rick White	T	Must define Ack_timeout value.	This value is either PHY dependent or based on the SIFS time.
5.2.9.	P. Brenner	E	The paragraph starting: "This policy..." should be immediately after the first paragraph	There is a paragraph in between that makes the whole last paragraph to be out of scope.
5.3	C. Heide	e	remove the extra "." at the end of the first paragraph.	
5.3	Jim Panian	E	Require all stations to be capable of participating in PCF data transfers during the contention-free period.	The last sentence of the introduction reads that "Nor, must all STA's be capable of participating in PCF data transfers." This implies that for power management, DTIMs cannot be scheduled during the contention-free period. Also, Beacons and ATIMs cannot be put out during the contention-free period.

5.3	bdobyns	T	Restriction that PCF cannot overlap coverage area with another PCF on same channel is fundamental flaw. Prohibits use of PCF with single-channel PHY, severely restricts use of PCF with N-channel PHY where N is small. PCF must be redesigned to permit the functionality of PCF to be delivered with single channel PHYs (or alternative way to deliver functionality). 1. superframe timing 2. cf period management 3. cf data delivery	
5.3	bdobyns	T	Append to paragraph 2 "The restriction on PCF operation for few-channel PHY need not preclude the overlap of Access Points for those PHY. It only precludes the PCF-style beacon behavior and a superframe delimited contention free service."	
5.3	C. Heide	t	specify what happens if there is a collision in the contention free period (by some STA which is hidden from part of the BSS).	since the PCF is build on the DCF, when the Point Coordinator backs off, what if a DCF STA gets control? Is there a gap in the CF period for the transaction to complete?
5.3	C. Heide	t	specify how a STA knows the start of a superframe.	the success of the contention free period is largely based on STAs knowing that it is in progress and setting their NAVs. If they can't do this, then partially hidden STAs will often see DIFS when there aren't any and destroy the CF burst.
5.3	C. Heide	t	second paragraph, second sentence, remove the clause "in a manner that results in destructive interference with frame transfer."	there is no manner of overlap that won't result in destructive interference, so why leave the excuse open?
5.3	C. Thomas Baumgartner	t	2nd paragraph is not acceptable to PAR requirements. Remove it. Can substitute discussion of mechanism decided on to handle overlapping point-coordinated BSA's.	There is no fully defined way to support time bounded services other than PCF. PAR says we must support time bounded services. Therefore can't restrict some (radio?) PHY's from using PCF. Not necessary to restrict any PHY. I propose a method to handle overlapping PCF BSA's (on same channel). If a STA hears PCF polls from an AP not it's own it tells its own AP of situation which causes the AP to change channels. I'm not sure how this works in FHSS since hop sequences eventually intersect but I'm sure we can quickly figure out how to handle this. Maybe the STA using FHSS concludes it is on same channel if it hears other AP polls twice within a number of hops. IR doesn't have overlap problem so not a problem that it has only one channel.

5.3	D. Johnson	T	Modify the PCF specifications to be consistent with a spectrum etiquette.	Operation in the US with unlicensed spectrum now, and likely in the future, will require conformance to an etiquette. The etiquette must operate on a power sense/timing basis only and must control fair access to any system conforming to it. This includes systems that do not necessarily conform to an interoperability standard. As a long-shot alternative, it may be possible to establish a regulatory arrangement in the US in which only one interoperability standard is permitted, much as HiperLAN in Europe. In this case the PCF could operate as specified, but this would require development of a very strong rationale. This change is highly recommended, but may not be possible in a convenient time schedule and is not a condition for changing the vote. A dialog should be set up with WINForum to determine what may be required to provide for a PCF function in the US in a future revision.	
5.3	Fischer, Mike.	T	replace text with: The Point Coordination Function (PCF) provides contention free services. It is an option for a station to become the Point Coordinator. All stations inherently obey the medium access rules of the PCF, because they are based on the DCF, with the Point Coordinator gaining priority access to the medium using a PIFS which is smaller than the DIFS used for the DCF access to the medium. The operating characteristics of the PCF are such that all stations are able to operate properly in the presence of a BSS where a PCF is in operation, and (if associated with a point coordinated BSS) receive asynchronous data frames send under PCF control. It is an option for a station to become the Point Coordinator, as well as to be able to respond to contentionÐfree polls by a Point Coordinator with CFÐdata transmissions. A station which is able to respond to contentionÐfree polls is referred to as CFÐAware, and may request to be polled by an active Point Coordinator during each superframe. When polled by the Point Coordinator, a CFÐAware station may send one data or CFÐdata frame to any destination (not just to the Point Coordinator). If the addressed recipient of a CFÐdata transmission is not CFÐAware, that station acknowledges the transmission using the normal DCF acknowledgement rules, and the Point Coordinator retains control of the medium using a PIFS duration before resuming CFÐtransfers.	update to match CFÐusage and subtype encoding that appear elsewhere in this draft

5.3	Geiger	T	The basic restriction is that a PCF can not overlap with another PCF on the same channel in a manner that results in destructive interference	This can not be guaranteed! Several things are missing from the standard to even come close to making this possible. First, there is no way to control a mobile unit's power. A 1000mW node can cross over several PCF functions with no control. Secondly, we specify no area restrictions in standard to limit range of PCF and associated nodes. This must be investigated further before it is either accepted or rejected in the Draft. I feel that this function must be removed.
5.3	Lewis	T	clarify limitations. If restricted to certain PHYs, state which PHYs. Need a clear mechanism to either control PCF contention or clearly define the set of conditions in which PCF is permitted or supported. If all conditions are not met, a station should not be permitted to intiate PCF.	
5.3	Rick White	T	¶ 2: If the use of PCF access is restricted to certain PHY type, these types must be defined. Otherwise this sentence should be removed.	Not defined
5.3.	Mahany	T	Add mechanism to control PCF contention in ESS or multiple independent BSS installations.	Second paragraph, it is not clear how PCF contention controlled in any PHY when multple BSS's are present simulataneously. IR is single channel , DS with limited channel set has high probability of overlap on a given frequency, or even in FH when occasional simultaneous usage of a given frequency will occur.
5.3.1	Geiger	E	Managed object SF_Length	My kids are manageable --- sometimes.... I think managed object is a better term for a manageable parameter. SF_Length is missing in the MIB definitions.
5.3.1	bdobyns	T	Define a Superframe Length Minimum and maximum (and put the min, max in the MIB).	
5.3.1	C. Heide	t	clarify the first paragraph and figure 5-15 to be consistent with figure 5-2.	The PCF runs over the DCF, they are not mutually exclusive.
5.3.1	C. Heide	t	explain where the length of CF parameter originates.	to preset their NAVs all STA must have the same parameter for length of superframe - from where to they get this and how is it assured they all have the same.
5.3.1	C. Thomas Baumgartner	t	Change 3rd sentence of 1st paragraph to " Within a given SF period, the PCF shall be active in the Contention Free Period, while the DCF is active all the time.	Figure 5.2 shows that DCF is basis for and always present with PCF. This is important to me because this was the major claim of superiority made for this protocol so we better not forget it.
5.3.1	C. Thomas Baumgartner	t	Change 1st sentence of 2nd paragraph to "The length of a SF is determined by the PC."	If manageable we need to say by whom
5.3.1	David Bagby	T	**1. Superframe Structure** Hum, the MAC group needs to discuss the super frame stuff again - I have heard it argued that the concept of the super frame is no longer applicable due to the way the CFA stuff has evolved. IF this is true then, the references to superframe need to be removed from the draft before sponsor ballot (refs in this section and others).[DB5]	See imbeded comments and annotations

5.3.1	Fischer, Mike.	T	The restriction of the superframe duration to not exceed the dwell time in an FHSS PHY renders the FHSS PHY essentially unable to support CF services of any kind with the typical dwell times listed in the current draft. This is unacceptable, so if the FHSS PHY is retained a means of permitting multipleÐdwell superframes must be found. Lengthening the dwell to several hundred milliseconds is not an appropriate solution, because of excessive delivery delays and variance that introduces to time bounded service when retries are necessary. I do not have a good solution to suggest, but urge that one be found or we will have to choose between no TBS or no contentionÐbased service with the FHSS PHY.	see column to right as well as other comments of mine regarding the dwell time limits (section 10.6.12)
5.3.1	Geiger	T	The hop dwell time is undefined	Define the period of time
5.3.1	Rick White	T	Superframe stretching must be removed.	There is not reason for it and it just complicates the synchronization of STAs. A STA should not transmit an Asynchronous frame if it and its ACK are not complete before the end of a superframe.
5.3.1	Wim Diepstraten	T	Superframe boundaries should be specified, such that they relate to the TSF. This allows stations to setup their NAV for a length of CF_Boundary, when they detect that a PCF is active in the BSS. The following is needed: "The target SF starting time will be when TSF mod SF_Length = 0." The text should also identify how the relevant parameters of SF_length and CF_Boundary are distributed. The prefered method is to put them in the Beacon.	The relevant parameters need to be put in the Beacon, rather then in the Association Response, because the latter will not work for overlapping stations of an other BSS.
5.3.1, et seq	Bob O'Hara	T	Delete the Superframe concept.	Superframe is a holdover from PCF TBS. Since PCF TBS is no longer supported, the superframe is no longer necessary. Because the PCF can gain priority access to the medium (through the use of PIFS) a superframe is not needed to support STAs in power saving modes either.
5.3.2	bdobyns	E	Drop the terminology "CF-up frames" and "CF-Down frames"	These frames are not distinguished from normal data frames, and the terminology is confusing. (there's no CF-UP in section 4).
5.3.2	C. Heide	e	first paragraph, third sentence, change to "Data frames sent from the PCF during the CF period to associated STA ...".	the PCF can sent data outside of the CF period also, these are not CF-Down frames.
5.3.2	Jeff Rackowitz	E	"...Contention Free Period may be accomplished using a bit in the header of subsequent Data frames... " What is the requirement?	
5.3.2	Jim Panian	E	Limit the use of the acronym "PC" to the first sentence of clause 5.3.2.	The text is confusing PC and PCF in this section and later.
5.3.2	C. Heide	t	first paragraph third sentence remove the word "associated".	Any STA can become a PC not just a AP according to previous paragraphs.
5.3.2	C. Heide	t	specify what bit in the header.	cannot find any such bit in section 4.
5.3.2	C. Thomas Baumgartner	t	Get these authors of this section together with the authors of Section 4 to determine where the CFP ACK bit should go in the frame header. There is a reserved bit in Frame Control field.	Wonderful improvement to the efficiency but not implemented in the frame format description of Section 4.
5.3.2	Fischer, Mike.	T	Add a statement that ÒThe PCF is not required to be located at the same station as the AP, but for most uses of contention free communication, any other configuration results in reduced throughput and increased transit delays for most frames.Ó	A sizeable percentage of BSS traffic is expected to have the AP as either transmitter or receiver. Hence the greatest gains from the lack of need for backoff and the piggybacking of acknowledgements comes when the AP station is TA of each CFÐdown frame and RA of each CFÐup frame.
5.3.2	Rick White	T	How is the "one special STA per BSS called the Point Coordinator" determined. This must be defined.	Not defined
5.3.2	Rick White	T	There are no longer any CF-Up and CF-Down frames. This section must be rewritten to reflect the currently defined Frame types and subtypes	
5.3.2.1	John Hayes	E	Change PCF to PC.	The coordinator is the PC, the function it provides is the PCF.
5.3.2.1	Bob O'Hara	T	Update to reflect current frame types in table 4-1	out of date

5.3.2.1	C. Thomas Baumgartner	t	Change last sentence to "Even if a DCF station do not set its NAV to the maximum CF-Period at the beginning of a SF for some reason, the shorter PIFS...needed by stations using DCF only."	Seems to me the setting of NAV is primary method of DCF deferal. Stations don't have a protocol called Contention period, the use DCF which results in their transmissions being in the Contention period.
5.3.2.1	Geiger	T	CF-End Frame needs to be defined	CF-End Frame is not defined anywhere
5.3.2.2	Bob O'Hara	E	replace "station with the PCF" with "Point Coordinator"	
5.3.2.2	Wim Diepstraten	E	Replace ....Superframe stretching. in the third sentence into .....collision.	
5.3.2.2	A. Bolea	T		It is not clear how a station knows when the SF is starting.
5.3.2.2	C. Heide	t	first paragraph, last sentence, replace "eliminates" with "minimizes".	because of hidden stations the possibility of a corrupted CF period cannot be eliminated.
5.3.2.2	C. Heide	t	first paragraph, first sentence should be "All non CF-aware STA shall preset their NAVs to the maximum ..."	If CF aware STAs set their NAVs, then they cannot transmit during the CF period.
5.3.2.2	C. Thomas Baumgartner	t	Must describe how STA knows the beginning of the SF so it can set its NAV.	This important operation not explained
5.3.2.2	C. Thomas Baumgartner	t	Add to end of 1st paragraph "Stations operating in PCF mode will ignore their NAV when they are directed to transmit by PC."	If everyone sets NAV at beginning of SF then no STA could transmit according to rules set up previously.
5.3.2.2	Fischer, Mike.	T	There appears to be considerable simplification to synchronizing the superframe with the beacon interval, especially in cases where the AP and PCF are colocated. Recommend adding a statement that ÒThere shall be an integral number of superframe intervals per beacon interval. The timing of these shall be synchronized such that the PIFS interval to gain medium access for one of these superframes immediately follows each beacon transmission.Ó	simplification, avoidance of severe delay in beacon generation when the nominal beacon interval occurs during a superframe (when the PCF and AP are not colocated) and the AP is not near the head of the polling list
5.3.2.2	Gegier	T	CF_Period parameter add to MIB table	CF_Period not listed in MIB table
5.3.2.2	Renfro	T		Describe how stations know superframe timing.
5.3.2.2	Rick White	T	How does a station know what the Maximum CF-period length is? This must be defined and must be a PCF distributed value.	
5.3.2.2	Tom T.	T	Replace the first sentence of the first paragraph with the following: The Duration field of the first Data frame sent by the PCF at the beginning of the Superframe shall have a value equal to the length of the contention-free period desired for this superframe by the PCF. Subsequent Data frames will have a duration field equal to the time remaining until the end of the Contention-Free period.	Throughout section 5.3 there is mention of the 'Start of the Superframe'. In this section each station is somehow supposed to know when this start is. There are two choices; use the duration field as mentioned to the left, or have a beacon or some kind of CF-Start frame at the beginning of each Superframe. (I would personally prefer a CF-Start frame, however the duration field solution builds on an existing capbility and has less impact on the standard)
5.3.2.2.	Fischerma:NAV Operation (Within the context of PCF operation)	T	First paragraph, missing reference: how do stations determine the beginning of the SF?	I cannot find a reference in D1 that indicates the mechanism for STATIONS to determine the beginning of the SuperFrame.
5.3.3	Joe Kubler	E	please make clear that PCF is always the AP.	
5.3.3.1	Bob O'Hara	E	update this section to reflect the deletion of the CF-ACK bit and addition of CF frame types	
5.3.3.1	Gegier	E	There is no reference to figure 5-17 in this section. Foreshortened Superframe? What is this?	Where is figure 5-17 explained? It is a good drawing but the detail for it is missing.
5.3.3.1	Joe Kubler	E	CF-Down should be defined as CF-DATA frame from PCF.	
5.3.3.1	Rick White	E	Figure 5-17 should be removed.	It is not referenced in the text.

5.3.3.1	Bob O'Hara	T	Add as the last sentence in the section: "CF frames shall not be retransmitted in the case of failure of acknowledgement."	The combination of priority access and retransmission without competition will lead to starvation of non-CF STAs in cases of noisy media or marginal transmission conditions.
5.3.3.1	C. Heide	t	third paragraph second sentence says "after a PIFS gap" - if this means without regard to busy medium it should say so.	if this is the case I want it pointed out so that I can object. If the response is there, but not being seen by the PCF, then transmitting blindly after a PIFS will corrupt it and the retransmission will never work and is just a waste of everybody's time.
5.3.3.1	C. Thomas Baumgartner	t	Second paragraph describes something that can't happen--a non CF aware station transmitting during the contention free period. I don't know solution but must be decided and changed.	But can a station that is not CF aware transmit during the CF period because it has set its NAV at the beginning of the SF period? I assume that all stations can ser NAV at beginning of SF since the description of that action didn't say otherwise. On thinking about this more I believe that pseudocode for sending of ACK and CTS in any circumstance includes ignoring NAV.
5.3.3.1	C. Thomas Baumgartner	t	change last sentence of 3rd paragraph to "...control and transmit the next frame after sensing that the medium is clear for a PIFS gap..."	Sentence doesn;t make it clear that PC is checking for medium busy during the PIFS gap.
5.3.3.1	Rick White	T	¶ 1: There is not a CF-Poll bit in the subtype field. CF-Polls are Async Data frame subtypes.	
5.3.3.1	Rick White	T	¶ 2: There is not a CF-Ack bit in the subtype field. CF-Acks are Async Data frame subtypes.	
5.3.3.1	Rick White	T	¶ 2: D2 must be for the same station if it id to be used for Acking the U1 frame. Otherwise the PCF will send a normal ACK.	
5.3.3.1	Rick White	T	¶ 4: In ¶3 it states that if a station is not PCF aware, it should responded with an ACK to a CF-Down frame. In ¶4 it states that a CF-Down frame need not be acked. Contradiction must be resolved.	

5.3.3.1.	Fischerma:PCF Transfers when the PCF station is Transmitter or Recipient	T	Last paragraph of section: Note that a station must at least respond with an acknowledgement for the preceeding CF-Down frame. The lack of any response to the CF-Down frame will be considered an error.	See first paragraph of this section: The PCF shall send (CF-Down) frames between the start of the CF–Period and the CF–End using the SIFS gap except in cases where a transmission by another station is expected by the PCF and an SIFS gap elapses without the receipt of the expected transmission. In such cases the PCF shall send the next (CF-Down) frame a PIFS gap after the end of the last transmission. A CF-Poll bit in the Subtype field of these frames will allow the stations to send their (CF-Up) data if any. Stations shall respond to the CF-Poll immediately when a frame is queued, by sending this frame after an SIFS gap. This results in a burst of Contention Free traffic; the CF-Burst. .... Last paragraph of this section: Note that a station need not respond when the station has no CF-Up traffic to send, and no acknowledgment is required to be returned for the preceding CF-Down frame. A responding CF-Up frame in these cases shall not be considered an error. This section has the following problem: How can the PCF know whether to expect a response or not if this is an option for the receiver? This forces the PCF to a choice of "always separate PCF transfers by PIFS, because there might be a response frame for any transfer." I.e.: If no acknowledgement is <u>required</u> to be returned, then the PC is likely to begin a new CF-down frame after SIFS, since the PCF is allowed to separate CF-down frames by SIFS when it does not <u>expect</u> a response. Now assume that the receiver of the first CF-down frame takes the "option" of generating a "responding CF-Up frame" in this case (even though it is NOT required), but the PCF is NOT expecting a response. The receiver will create a collision with the PC because it has taken this option. Therefore, the "option" should be stricken from the specification, and the generation of a responding CF-Up frame when the PCF is not expecting it should be considered an ERROR and illegal. Basically, the problem here is that there needs to be explicit requirements on the part of the CF-aware station, such that the PCF can determine whether or not a response is forthcoming. If the PCF cannot make this determination, then the rule earlier in this section needs to be changed, such that <u>all</u> CF-down traffic for which a response cannot be predicted should be separated by PIFS instead of by SIFS. Such a delineation is not currently clear.
5.3.3.2	C. Heide	e	first line, insert space in "thePCF".	
5.3.3.2	C. Heide	e	first paragraph twice refers to "an SIFS" instead of "a SIFS".	

5.3.3.2	C. Heide	t	station to station transfer in the CF period should not be allowed	(1) if the PC can't hear the destination station the transfer will never work and will corrupt forthcoming transmissions from the PC. (2) the length limit of the CF period cannot be guaranteed because the PC has no control over how long these two STA seize the medium.
5.3.3.2	C. Thomas Baumgartner	t	Add after 2nd sentence "To allow this transaction the PCF, when it receives a data frame no directed to it, waits PIFS instead of SIFS.	This action is implied but better to explicitly state it.
5.3.3.2	Mahany	T	Change Max MPDU Length (Figure 5-17) to MSDU Length	Must allow for full MSDU, with fragmentation. Also see Mahany comment at 5.3.4.1
5.3.3.2	Renfro	T		States that PCF will resume CF-Down transmissions after SIFS period after ACK. What if message being acked contains additional fragments? Either PCF must be assured of hearing message or PCF must wait PIFS after ACK to begin transmission.
5.3.4	C. Heide	t	define TBS that "may have multiple service levels."	what does that mean?
5.3.4	Fischer, Mike.	T	Replace this text with ÒThe PCF provides an alternate mechanism to access the WM. Within this contention free medium access, both asynchronous and time bounded services can be provided.Ó	clarify that CF is not a service but a medium access modality
5.3.4	Rick White	T	If Contention Free Time Bounded Services "may have multiple service levels", they must be defined.	Not defined.
5.3.4.1	Bob O'Hara	E	replace "for" with "to" in the last paragraph	
5.3.4.1	Tom T.	E	Change 'to these services' on second line to: 'to Contention Free services'. Change 'during the Superframe period' on last line of first paragraph to: 'during the Contention Period'.	
5.3.4.1	Fischer, Mike.	T	Change equation for CF_Boundary to be Max. Fragmented MPDU with RTS/CTS and ACKÓ	leave room for a full contention-based frame per superframe
5.3.4.1	Geiger	T	How does BACKoff operate during the end of the superframe? CF_Boundary missing from MIB. Max. Async MPDU is undefined and missing from MIB.	The amount of time allowed for the contention period, one max size MPDU is silly. I believe that the contention free period should take no more than 1/2 of the superframe, especially if it is intended for real-time traffic. This allows bursty traffic to get through at a reasonable rate without moving up to priority type of traffic. The priority traffic doesn't have any defined mechanism for congestion control. This is unacceptable.
5.3.4.1	Geiger	T	Some people would like to see the dwell time in FHSS PHY not exceed more than a few max packet lengths. This is to avoid interference from microwave ovens etc. We might consider that a Superframe extend over several hop & hop dwell periods rather than size it to a single dwell time	One consideration for Superframes to contain several hops & hop dwell periods is that one might be able to scale hop times dynamically without impacting the Superframe.
5.3.4.1	Greg Smith	T	The contention free period shall be limited to 50% of the SuperFrame	Having a mechanism in the standard that allows one implementatior aware) to shut down async only stations to one packet per SuperFra could be considered predatory. After all in an ISM band the CF peri going to have to cope with other outside interference, why not async traffic.
5.3.4.1	Mahany	T	Revise so that superframe allows at least one fragmented Asynchronous MSDU may be transmitted per superframe using RTS, CTS mechanism, , with fragment length set to aMPDU_Minimum. Alternatively setting fragment length to aMPDU_Current_Maximum would be acceptable.	It is not clear whether the provision for one max. Asynchronous MPDU allows for fragmentation, RTS, CTS collision avoidance, etc. This must be made explicit, as it may preclude use of some access mechanisms (or fragmentation) if PCF is used, or it may force adaptive algorithms to establish SF length.
5.3.4.1	Renfro	T		Must allow time for max MPDU and Max contention window. If large amount of contention traffic, superframe will continually stretch to the point where there will be insufficient time for contention free service in a particular superframe.

5.3.4.1	Rick White	T	Must define a MIB Value for length of the Contention Free period.	Not defined.
5.3.4.1	Rick White	T	¶ 2: Superframe stretching must be removed.	There is not reason for it and it just complicates the synchronization of STAs. A STA should not transmit an Asynchronous frame if it and its ACK are not complete before the end of a superframe.
5.3.4.1.	Fischerma:Contention Free Length Limit	T	Size of superframe should be revisited, based upon objective.	The reasoning given for the chosen value of the limit is nonsensical. The requirement needed in order to guarantee the time necessary for "at least one maximum size Asynchronous MPDU" to be transmitted during the contention portion of the superframe is infinite. This is because in a heavily loaded network, it is possible (although not likely) for collisions to consume the entire contention portion of the SF. The choice of the size of one maximum size Asynchronous MPDU increases the probability of transmission of a contention period frame, but it does not guarantee it. Once a frame begins, the medium should be sensed busy by the PCF, and therefore, SF stretching should result to allow any size asynchronous frame to be transmitted. If the point is to avoid SF stretching, then this method might help to minimize SF stretching, but there is still the possibility of a first contention frame (or a combination of frames) using 99% of the contention period, and the next asynchronous frame being a maximum size MPDU then stretching the SF into the next CF period for the maximum possible amount of time anyway. The limit should be set based upon the 95% confidence interval for collision resolution given a "large" number of contending nodes all attempting to gain control of the network immediately following the end of the CF period PLUS some period of time (like the maximum length asymchronous MPDU) in order to attempt to guard against excessive SF stretching. Probability of SF stretching is very likely in any case anyway.
5.3.4.2	Bob O'Hara	E	remove numbers from beginning of paragraphs	
5.3.4.2	A. Bolea	T		References to CF-Poll, CF-ACK Bits need to be corrected using new frame types. This applies to other sections also.
5.3.4.2	Bob O'Hara	T	replace the second sentence of paragraph two with: "A CF-aware station shall acknowledge receipt of each Asynchronous Data frame of the CF-Poll subtypes from the PCF using Data frames of the CF-ACK subtypes, sent after an SIFS interval. A CF-aware stationshall acknowledge receipt of all other Asynchronous Data frames using ACK Control frames sent after an SIFS interval.	Update to reflect new frame subtypes.
5.3.4.2	C. Heide	t	rule 1 - previous sections say that RTS/CTS use is controlled by the RTS_Threshold parameter. Clarify how this rule is broken in the CF period.	contradicts section 5.2.7
5.3.4.2	C. Heide	t	rule 3 - correct the last sentence, a STA is allowed to respond or not respond.	contradicts 5.3.3.1
5.3.4.2	C. Thomas Baumgartner	t	Change to "1. Only Data frames, resulting ACK frames (if any), RTS frames, and resulting CTS frames shall be sent..."	Sentence was incorrect. RTS/CTS is allowed to enhance reliability of CF transmissions against hidden nodes and to aid in resolving overlapping BSA contention.

5.3.4.2	Fischer, Mike.	T	There appears to be no reason for prohibiting management frames during the CF period. Suggest deleting usage rule #1.	In general, management frames can be send anywhere data frames are allowed, hence they should be permitted during the CF period.
5.3.4.2	Rick White	T	Contention Free usage rules must be rewritten in order to eliminate the reference to bits and instead reference frame subtypes.	There are no longer CF bits in the control field of the MAC header. They are different Frame types.
5.3.4.2	Wim Diepstraten	T	All management frames should also be allowed to be send during the CF-period.	The limitation is not necessary, and does greatly increase the Beacon timing complexity, because the SF-interval and Beacon interval can not be alligned.
5.3.5	C. Heide	e	second paragraph, first sentence, spelling error "receive"	
5.3.5	Joe Kubler	E	2nd para, "recive CF-Polls" should read "receive CF-Polls"	
5.3.5	Renfro	E		Add ACFS to acronym table.
5.3.5	bdobyns	T	Permit ACFS to be initiated in a non-PCF environment with a *Start Connection Request*	This allows few-channel PHY which cannot support PCF functions because of overlap restrictions to still support a contention-free service which may have a better Qos than the regular Async service.
5.3.5	Bob O'Hara	T	Update paragraphs to reflect new subtypes in table 4-1	Out of date.
5.3.5	Geiger	T	Asynchronous traffic is characterized by its bursty, connectionless nature. The ACFS allows	This paragraph is a bunch of bull. The difference between data transfers in the contention free period versus the contention period is the QoS. The contention free period allows a connection oriented service to be established with some QoS associated with the access to the media and predictable results when the media is busy or no more contention free period is available. The contention period provides no QoS. No bandwidth can be reserved or guaranteed nor can delay or congestion be managed. Access to the contention free period is managed by the PFC. Access to the media during the contention period is asynchronous in the sense that the point in which a stations grabs the media is not predictable by any other station in WLAN. Accesses in the CF period are predictable by the PCF and other station in the WLAN. Talking about a ACFS procedure is silly. The discussion here should be about connection oriented or connectionless services, not ACFs what every that means.
5.3.5	Geiger	T	The polling list is a logical construct,... This list has to be more than this. If the intent of the CF period is to provide better QoS than the contention period, how the QoS is implemented is a function of the MAC.	If the MAC is going to control the access to the media, and part of that access involves some QoS parameters, then how the MAC administrates the QoS of the polling list better be in the standard. Managing the CF polls and who needs service versus who doesn't during each frame must be a function of the MAC, not some higher layer.
5.3.5	Rick White	T	Must define how a station gets on the "polling list".	
5.3.5.1	C. Heide	e	the last sentence of the last paragraph is unintelligible	can't correct it because I don't understand it.
5.3.5.1	Miceli	E	"The PCF is not required to do this, and in certain cases, such as a (CF-Down) frame that acknowledges a (CF-Up) frame less than one MPDU duration from the CF-boundary, the CF-Poll bit must not be set."	existing text is confusing
5.3.5.1	Bob O'Hara	T	Update paragraphs to reflect new subtypes in table 4-1	Out of date.
5.3.5.1	Bob O'Hara	T	Paragraph two is difficult to understand and must be rewritten	Ambiguous
5.3.5.1	C. Heide	t	clarify what the PC does if the superframe has been delayed s long there is no longer time to send at least one Data frame.	conflict between requirement to send at least one data frame and to restrict the maximum length of a superframe and allow superframe start delay.
5.3.5.1	Rick White	T	Must define whether the PCF can change the CF_Boundary based on the amount of CF traffic expected.	Not defined.
5.3.5.1	Rick White	T	Must define how the PCF works through the "polling list". If it is not completed during a CF period, does the PCF start over the next period or pick up where it left off?	Not defined.

5.3.5.1	Rick White	T	Must rewrite to reference Data frame subtypes, not bits in the header.	There are no longer CF bits in the control field of the MAC header. They are different Frame types.
5.3.5.2	Rick White	T	How a STA gets on the "polling list" must be inside the scope of the standard and must be defined. A mechanism must be defined to allow a station to be added to the "polling list".	
5.3.6	Wim Diepstraten	E	Change the sentence to: "The contention free TBS management frames are used in the following way.	The management frames listed are for CF-TBS only. This should be made clear in the text.
5.3.6	C. Heide	t	specify how a connection request is denied.	how is a connection request denied - sending an End Connection in response to a request Connection? This section doesn't say.
5.3.6	Rick White	T	Must define all frames and the content of each of the fields.	This whole section must be rewritten with more detailed information on how the frames are used and what happens when a frame is received. Since they are management frames, they are not passed up to or received from the LLC.
5.3.6.1	Mahany	E	"MAC User (of an STA)should be "CF Aware STA"	Clarity
5.3.6.1	Bob O'Hara	T	Define "Start Connection Request"	There is no "Start Connection Request" defined in the MAC service interface.
5.3.6.1	Bob O'Hara	T	Define "Start Connection Indication"	There is no "Start Connection Indication" defined in the MAC service interface.
5.3.6.2	C. Heide	e	add "." to end of first sentence	
5.3.6.2	Mahany	E	"MAC User (of an AP)should be "CF Aware AP"	Clarity
5.3.6.2	Mahany	E	Replace N.B. with plain English.	Clarity
5.3.6.2	Renfro	E		What is N.B. AP?
5.3.6.2	Tim Phipps	E	*Remove: "NB AP and STA start connection ... them".*	There is no need to distinguish the frames. The "ToAP" bit no longer exists.
5.3.6.2	Bob O'Hara	T	Define "Start Connection Request"	There is no "Start Connection Request" defined in the MAC service interface.
5.3.6.2	Bob O'Hara	T	replace "N.B." with appropriate standard language and functional description.	Proper standard language required
5.3.6.2.	Fischerma:AP Start Connection Request	T	Last pargraph of this section: N.B. AP and STA Start Connection Request frames are the same type, using the "To AP" bit to distinguish them.	Let's be definite about the type designation - the two type fields are identical or they are not - the original text used the term "can" in the sense of they might be if you want them to be...
5.3.6.3	C. Heide	e	first sentence, remove the word "Start", and the long quote mark.	
5.3.6.3	Mahany	E	MAC does not reply with frames. Replace with "Point Coordinator"	Clarity
5.3.6.3	C. Heide	t	last sentence, change the second "connection" to "STA".	the STA is what gets added to the poll list, not the connection.
5.3.6.3	Tom T.	T	Change Grant Connection frame type to Connection Response. Change in first line: 'MAC may reply' with 'MAC shall reply'. Delete first sentence of second paragraph. Replace third paragraph with: The connection may be granted or denied by the AP and shall indicate this using the Status Value and Error Indicator elements.	It is better to make these exchanges more deterministic. Getting no response at all gives the higher layers no information about what's happening, therefore a negative response should be used.
5.4	Okada	E Approve	PRNG is not defined	

5.4	Bill Huhn	T	Privacy must be included as an 802.11 standard.	Customers will demand some sort of security from standard 802.11 product and will not want to add this on. Lack of integral privacy will slow market adoption of 802.11 standard based product.
5.4	Fischer, Mike.	T MAJOR ISSUE	This section needs moderate update to accommodate the WEP recommendations from the January, 1995 interim meeting. Document 95/15 contains recommended replacement text and drawings to accomplish this. An alternative ICV algorithm is presented therein, which has certain advantages (also stated therein), but CRC32 is also an acceptable, if suboptimal, ICV technique. The PRNG algorithm remains unspecified in document 95/15. I found the proposed RC(4) algorithm from the original WEP proposal to be acceptable. However, the changes discussed in document 95/15 deal with aspects of WEP that are orthogonal to the PRNG algorithm details, and are applicable to any PRNG algorithm that uses (KEY+IV) up to 64 bits and uses XOR to encipher the L_SDU to yield and SDE_SDU.	The objective is a WEP that is practical and efficient to implement as a generally available privacy mechanism, adequate for basic privacy and able to work, or be disabled, in environments which want more comprehensive security facilities, and operate and 802.10 SDE layer above the MAC. This avoids the complexity of potentially needing a full 802.10 layer within the MAC, which would discourage implementation and use of the basic privacy facilities.
5.4	Lee Hamilton	T	Privacy must be defined as a part of 802.11 and not require 802.10.	This is a key to market acceptance of a wireless LAN. Customers will not accept product that requires additional additions to implement privacy.
5.4	Lee Hamilton	T	The pseudo random generator must be defined.	The privacy algorithm is not implementable without the PRNG.
5.4	Rick White	T	It is not clear from this section that the WEP algorithm is implemented as part of 802.10.	
5.4	Rick White	T	The WEP must be part of the MAC. It should not require any implementation of 802.10. It should be independent of 802.10. If more security is required, the WEP must have the capability of being disabled and 802.10 security used above the MAC.	Customers will require privacy on their WLANs. They will not what to be required to use another standard to implement it.
5.4	Rick White	T	The pseudo random number generator must be defined.	Not defined. I vote for the one that was supplied at the Nov. 94 Meeting.
5.4	Scaldeferri	T	Text will be provided at the March Meeting	This section on Wired Equivalent Privacy (WEP) needed to be expanded to show how it fits the IEEE 802.10 standard. It should include at least two more figures and associated text showing the WEP's SDE structure and the SDE local management flow. It should also state and describe that WEP will provide for implicit authentication as described in 802.10
5.4, 2.4.3.2	Jim Panian	T	For conformance, support for the WEP privacy algorithm (or other standardized privacy algorithm) must be static (must be implemented). The actual use of the WEP privacy scheme may be dynamic (may not be used on every association).	Why isn't a standard privacy algorithm specified? The lack of a standard specified privacy algorithm will hinder interoperability.
5.4.1	C. Heide	t	second paragraph, remove the last two sentences.	contradicts section 2.4.3.1 which says that authentication is mandatory.
5.4.1	C. Thomas Baumgartner	t	Change 2nd sentence to "WEP is an 802.11 option to provide a data confidentiality algorithm. Change last 2 sentences to "Running an 802.11 network with privacy but without authentication leaves the system open to security threats."	I don't remember a vote where P802.11 committee as a whole recommended against privacy without authentication. The new sentences are statements of fact without the emotional baggage.
5.4.1	Wim Diepstraten	T	Delete the last two sentences should be deleted.	At the 802.11 MAC level we can assume implicit authentication whenever privacy is used, and the frame decripts succesfully as can be detected by the ICV check mechanism.
5.4.2	bdobyns	E	is 'exportability' paragraph appropriate for an international standard?	
5.4.2	Bob O'Hara	E	Replace "exportability" with "Export"	consitent usage
5.4.2	David Bagby	E	**Exportability:** Every effort has been made to design the WEP system operation so as to maximize the	See imbeded comments and annotations

5.4.2	C. Thomas Baumgartner	t	Delete last sentence.	The exportability of WEP is only one reason that it is an option. IR doesn't need it as much as RF for instance could be another reason. The desire of customers not to deal with key management is another good reason.
5.4.3	Bob O'Hara	E	delete "proposed in this submission"	
5.4.3	C. Heide	e	paragraph following figure 5-21, remove "proposed in this submission".	
5.4.3	Bob O'Hara	T	add "via a secure side channel" after "needs to be communicated between stations"	State explicitly what is only implied.
5.4.3	C. Heide	t	define variable P .	second paragraph following figure 5-22 refers to variable P which is not defined.
5.4.3	Renfro	T		k has been defined as both the key and the key sequence. Need to clarify.
5.4.3	Renfro	T		Since the encrypted message now includes IV and ICV, does this mean that the maximum MSDU size of 2034 bytes is actually increased?
5.4.3	Tim Phipps	T	*A number of details in this section do not comply with the requirements of 802.10.* *Replace figure 5-22 with:* 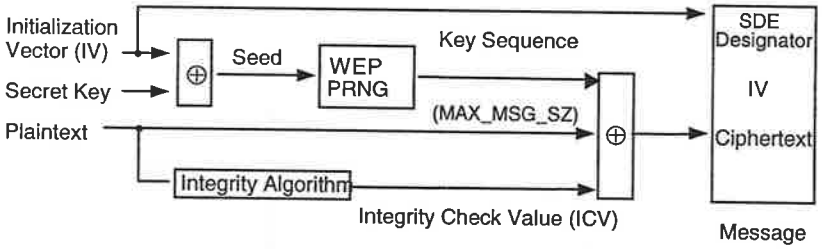 *Replace:* "The output of the process .. ICV." *with:* "The output of the process is a **message** containing the ciphertext, the IV, and other protocol fields required by 802.10". *Replace:* "The {IV, MSDU, ICV} triplet forms" *with* "The {IV, MSDU, ICV} triplet, and any other protocol fields required by 802.10 (such as the SDE desginator, which is a special LSAP indicator), form"	802.10 requires that if a clear header is present, it must include the SDE designator. 802.10 requires that the ICV be within the protected part of the message. 802.10 requires the SDE designator. In addition, a full implementation of 802.10 allows additional protocol fields to be present within the clear header.

5.4.3	Wim Diepstraten	T	Clearly specify in the text and pictures, how the Secret Key, and the Initialization Vector (IV) are combined to form the seed of the WEP PRNG. To support implicit authentication, it should be specified that a frame that is received with a correct FCS, but where the ICV does not check shall not be send to the LLC. The WEP integrity check algorithm should be changed. It should be an algorithm that can easily be applied on a per fragment basis, such that per fragment ICV check values can easily be combined to the MSDU level ICV. If privacy is optional then there should be an indication in the MAC header as to whether privacy has been applied to this frame. Privacy should only apply to the MSDU, not to the MAC Header, nor to Management and Control frames. 802.11 should specify that by default an ESS-wide key is to be used. 802.11 should specify an element that can optionally be used in a Beacon, for the purpose of IV distribution management.	The picture suggest an EXOR, while the operation is apparently concatenation. Only frames that are using the same key for encription and decription are valid for forwarding to the LLC. However the frame should not be discarded by the MAC itself, because it does contain relevent information like the Duration and PM bits. The ICV check algorithm is unnessesary complex. It should be noted that the same CRC HW implementation cannot be reused for ICV check purposes. A more simple reversible algorithm should be specified that can be applied in fragments or parts thereof. For efficient implementation within the MAC this indication is needed. This is to facilitate a simple default mechanism, in which only one key is used network wide. This should allow efficient implementations in which the AP can manage the prefered use of an IV.
5.4.3.2	Bob O'Hara	T	Update paragraphs three and four to reflect new subtypes in table 4-1	Out of date.
5.4.4	C. Heide	e	remove second sentence.	
5.4.4	A. Bolea	T		The PRNG needs to be specified.
5.4.4	bdobyns	T	must specify the PRNG algorithm	maybe you can use the one from 5.2.5 (heh, heh)
5.4.4	Bob O'Hara	T	an algorithm must be specified to provide the required security.	The function is described but not defined.
5.4.4	C. Thomas Baumgartner	t	Delete last sentence. Replace with "Until the next version of this specification details the PRNG this feature can't be implemented in 802.11 compliant networks for interoperability reasons."	Have to say this unless there is a method for WEP capable STA's to keep track of each STA's capability in ESS and send in clear when the destination is not WEP capable.

5.4.4	David Bagby	T	**2.WEP Algorithm Specification** The specific PRNG algorithm is unspecified at present. Reviewers of this draft are encouraged to comment on appropriate PRNG algorithms for adoption by 802.11. The PRNG must be specified to enable compatible implementations of WEP. This reviewer has concluded that for several reasons RC4 from RSA is the appropriate alg to choose. The reasons include effciency of implementation and the enhacned chances for export offered by the unique status of RC4. This reviewer can not vote to forward the Standard for sponsor ballot until the PRNG alg is specified.	See imbeded comments and annotations
5.4.4	Dean Kawaguchi	T	WEP Algorithm Specification	Currently TBD
5.4.4	Geiger	T	WEP algorithm is unspecified.	Specify
5.4.4	Jon Rosdahl	T	See page 9 of Document IEEE P802.11-94/249 for exact algorithm	I think that the algorithm described in Doc IEEE P802.11-94/249 would make wireless LANs actually more secure than the wired equivalant. This seems to be a very good reason to use it if possible.
5.4.4	Renfro	T		Need to be careful in defining WEP algorithm. If products containing WEP are not exportable, a standard defining WEP may not be exportable either.
5.4.4	**Siep**	**T**	**WEP Algorithm Specification[an export-approved algorithm must be specified]**	**This reflects the discussions on Encipherment held in the January MAC meeting in San Jose. The minimum criteria is:** • **exportable from the US** • **importable to France** **If this proves to be impossible, then all references to encription and its requirement in 802.11 should be removed.**
5.4.4	Stuart Kerry	T	WEP Algorithm Specification	Currently TBD
5.5	CHRIS ZEGELIN		SIMPLE STATEMENT THAT ONCE THE POLL HAS BEEN TRANSMITTED THAT PSP STATIONS RECEIVE THE MSDU JUST LIKE A CAM STATION	HOW FRAGMENTATION WORKS FOR 'PSP' STATIONS IS NOT ADEQUATELY SPECIFIED.
5.5	Bob O'Hara	E	replace "needs" with "is" in the third paragraph	
5.5	Bob O'Hara	E	add "for an MSDU of 1500 octets" to the end of paragraph five	
5.5	Bob O'Hara	E	change all "bytes" to "octets" in figure 5-24	
5.5	Bob O'Hara	E	replace "must" with "shall" in paragraph seven	Proper standard language
5.5	Bob O'Hara	E	change all "bytes" to "octets" in figure 5-25	
5.5	Bob O'Hara	E	update to reflect new sequence control semantics.	
5.5	C. Heide	e	last paragraph last sentence, replace "than" with "then"	

5.5	Geiger	E	Whenever possible, the size of the payload of a fragment shall be some fixed number of bytes	Good, I would hate to see it be some variable number of sheep. What does this mean?
5.5	Renfro	E	In 4th Paragraph change '... following two ...' to '... following three ...'. Add c) aFrag_Payload. Update references to MSDU ID and fragment ID to reflect Sequence Control Field	
5.5	Rick White	E	¶ 10: Change MSDU ID to Sequence Number.	MSDU ID no longer used.
5.5	Tim Phipps	E	When data needs to be transmitted, the number of octets in the payload of the fragment shall be determined based on the time at which the fragment is to be transmitted for the first time. Once a fragment is transmitted for the first time, its contents shall be fixed until **the MSDU** is successfully delivered to **the DS or destination station.**	The specification said that the fragment size must be kept constant until the MPDU reaches the destination. Fragmentation is applied within a BSS, different BSSs will chose different fragment sizes. Therefore, when more than one 'radio 'hop' is used, fragment size cannot be fixed until the ultimate "destination" is reached.
5.5	Tom T.	E	Change 'MSDU ID' to: 'Dialog Token' in third last paragraph.	
5.5	A. Bolea	T		References to MPDU ID need to be replaced with Sequence Control. Last Fragment bit is now in Frame Control Field and not in Fragment Number.
5.5	bdobyns	T	An implementation whose PHY MIB parameter aMPDU_Minimum is greater than 2304 plus MAC Header may choose to not implement fragmentation on either transmit or receive.	
5.5	Bob O'Hara	T	insert "assembled" between "is" and "to be" in the fouth paragraph	further clarification of when fragmentation takes place.
5.5	Bob O'Hara	T	in paragraph seven change the second sentence to be "In this case, the station shall wait until after the well boundary to create..."	Better clarity
5.5	Bob O'Hara	T	delete paragraph eleven	Unecesary complexity to squeeze, on average, half a frame into each hop period.
5.5	Bob O'Hara	T	Define all atributes in the MIB in section 7	These attributes are not defined.
5.5	C. Heide	t	remove references to MSDU ID.	MSDU ID undefined
5.5	C. Heide	t	define aTransmit_MSDU_Timer attribute	section 4 frame descriptions do not define this.
5.5	C. Thomas Baumgartner	t	Authors of this section need to get with authors of frame format section and decide where the MSDU ID will be.	No MSDU ID in section 4 frame format description

5.5	David Bagby	T	**3. Fragmentation** combine this section with sec 5.1.5 so frag info all in one place[DB8] After due consideration, and recognizing that stations are explicitly not required to attempt to fit fragments to remaining dewell times fir FH PHYs, and considering that the increase in band width utilization involved is very slight, I conclude that the complexity of attempting to match fragment size to remaining dwell time does not justify the effort involved. Even as an option, I don't believe we should retain this feature as the draft is already the most complex MAC ever defined. This is an area were we should increase the odds of interoperability and simplicity over functionality. Therefore, I vote against sponsor ballot until this feature is removed. If this modification is adopted, I shall volunteer to edit sections 1.1.4 and 5.5 to make the needed wording changes. I have not provided exact text here as word does not allow recursive annotations and that change would obscure other comments I have made in the same sections.[DB9] The MAC may~~will~~ fragment and reassemble MSDUs. The fragmentation and reassembly mechanisms allows for fragments to be retransmitted.	See imbeded comments and annotations

5.5	David Bagby continuation	T	one the consequences of providing fragmentation at the MAC layer is that a station must contain have MSDU buffering to cover ((max MSDU size + MAC overhead) * number of ad-hoc stas one wishes to communicate with simultaneously). This is true for both infrastructure and IBSS operation. To provide a minimal level of interoperability, a minimal number for simulations station support must be specified. this is on the order of 2k+ per simultaneous station and may not be an insignificant implementation cost. Once the number of different MSDUs being received exceeds the available buffering, there will be a failure condition. It is my assumption that the way this failure will manifest itself is that new MSDUs will not be received and therefore not acked, eventually resulting in retransmission (hopefully when the number of simultaneous MSDUs being received at the destination is less). To guarantee some level of avoidance of this problem, we must specify a minimally supported number of simultaneous MSDU receptions. to do this the following sentence should be added. I have chosen 6 MSDUs as it adds up to a bit less than a common memory increment. All Stations shall support the simultaneous reception of a minimum of 6 MSDUs. The fragmentation mechanism design accounts for the characteristics of FH PHYs. For the purposes of this description a 'dwell time' will refer to the duration of time spent on a single frequency in a FH system. Therefore in a FH PHY, the PHY will hop to the next frequency in the hop sequence at the end of the current dwell time. For other systems a 'dwell time' will refer to the period of time spanning from the start of transmission of a TIM until just before the start of transmission of the next TIM.	

5.5	David Bagby continuation	T	Whenever possible, the size of the payload of a fragment shall be some fixed number of octets. This is denoted by aFragment_Threshold~~Payload~~. aFragment_Payload equals ~~aFragmentation_Threshold minus MAC Header minus CRC~~. The payload of a fragment can~~shall~~ never be larger than (aFragmentation_Threshold - MAC Header Length - CRC Length)~~aFragment_Payload~~. However, the size of the payload may be less than this~~aFragment_Payload~~. When data needs to be transmitted, the number of octets in the payload of the fragment shall be determined based on the time at which the fragment is to be transmitted for the first time. Once a fragment is transmitted for the first time, its contents shall be fixed until it is successfully delivered to the destination station. The number of data octets in the payload of a fragment shall depend on the values of the following two variables at the instant the fragment is to be transmitted for the first time: a) The time remaining in the current dwell time. b) The number of octets in the MSDU that have not yet been transmitted for the first time. Since the control of the channel will be lost at a dwell time boundary and the station will have to contend for the channel after the dwell boundary, it is required that the acknowledgment of a fragment be transmitted before the stations cross the dwell time boundary. Hence, if there is not enough time remaining in the dwell time to transmit a fragment with an aFragment_Payload payload, the number of octets in the payload may be reduced to the maximum number of octets that will allow the fragment plus the MAC acknowledgment to fit within the time remaining in the dwell time. This is shown in Figure 5-24.	

5.5	David Bagby continuation	T	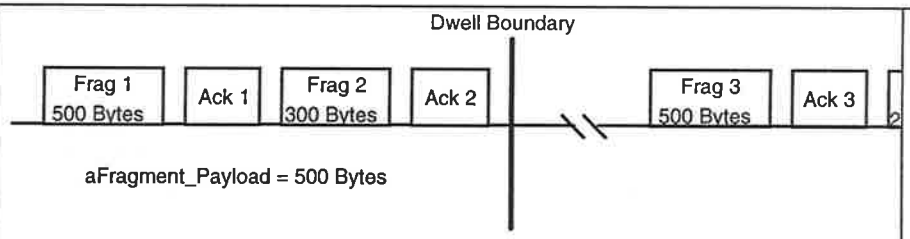 **Figure 5-24: Fragmentation Near a Dwell Boundary** Referring to Figure 5-24, a example 1500 octet MSDU is fragmented into four fragments with aFragment_Payload set at 500 octets. There is enough time left in the dwell to send two fragments, one of 500 octets and a second of 300 octets. After the dwell boundary, the rest of the MSDU is sent, one 500 octet fragment and one 200 octet fragment. A station may elect not to adjust the size of the fragment~~payload~~ when approaching a dwell boundary. In this case, the station shall wait until the next dwell time to create and transmit a fragment, ~~with a aFragment_Payload octet payload (provided there are at least aFragment_Payload more octets remaining in the MSDU). A~~ station must be capable of receiving fragments of ~~varying~~ sizes varying between aMin_Full_MPDU and aMax_Full_MPDU for a single MSDU. If a fragment requires retransmission, its contents and length shall remain fixed for the lifetime of the MSDU ~~at that station. In other words, A~~after a fragment is transmitted once, the contents and~~or~~ length of that fragment shall~~are~~ not ~~allowed to change~~fluctuate to accommodate ~~the~~ dwell time boundaries. For example; Let the fragmentation set refer to the contents and length of each of the fragments that make up the MSDU. The fragmentation set is created at a station as soon as the fragments are attempted for the first time. The fragmentation set remains fixed for the lifetime of the packet at the transmitting station. This is shown in Figure 5-25.	

5.5	David Bagby continuation	T	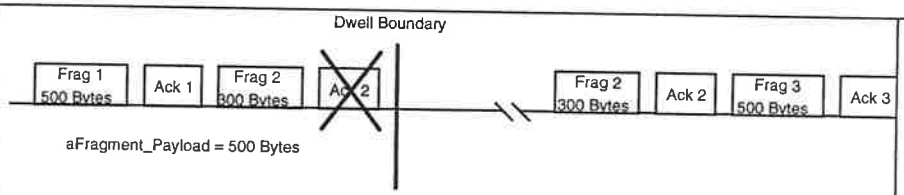 **Figure 5-25: Fragmented MSDU with missed ACK Near a Dwell Boundary** In the example shown in Figure 5-25, the same 1500 octet MSDU is fragmented at the same point in the dwell time as in Figure 5-24 but the ACK for the second fragment is missed. After the dwell boundary, the fragment is retransmitted and the fragment size remains 300 octets. Each fragment will contain a MSDU ID and fragment ID. When a station is transmitting a MSDU, the MSDU ID will remain the same for a given MSDU and the fragments will be in order of lowest ID to highest ID. The fragment ID also contains a bit that indicates the last fragment of the MSDU. If, when retransmitting a fragment, there is not enough time remaining in the dwell time to allow transmission of the fragment plus the acknowledgment, the station shall wait until the start of the next dwell time before retransmitting that fragment. The source station will maintain a aTransmit_MSDU_Timer attribute for each MSDU being transmitted. There is also an attribute, aMax_Transmit_MSDU_Lifetime, that specifies the maximum amount of time allowed to transmit a MSDU. The aTransmit_MSDU_Timer starts on the attempt to transmit the first fragment of the MSDU. If aTransmit_MSDU_Timer exceeds aMax_Transmit_MSDU_Lifetime than all remaining fragments are discarded by the source station and no attempt is made to complete transmission of the MSDU.	
5.5	John Hayes	T	TBD	This section does not address how to fragment broadcast and multicast frames.
5.5	Mahany	T	First Paragraph: Correct Definition of Dwell Time to be Applicable to All PHY's per 5.3.1	Statement that interval between TIM's of FH time on frequency defines dwell time is incorrect if PCF is used.
5.5	Paul Pirillo	T	Timing diagrams and/or text should be modified to show that aMax_Transmit_MSDU_Lifetime is actually the "dwell time." Define the relationship between aMax_Transmit_MSDU_Lifetime and the SF Period defined in section 5.3. Or if there is no relationship state so.	I am unclear as to what parameters define "dwell time." May want to consider modifications to section 5.2.6.5 as well, to improve clarity. I also am unclear as to how the PCF environment affects fragmented MSDUs.

5.5	Paul Pirillo	T	Timing diagrams and/or text should be modified to show that aMax_Transmit_MSDU_Lifetime is actually the "dwell time." Define the relationship between aMax_Transmit_MSDU_Lifetime and the SF Period defined in section 5.3. Or if there is no relationship state so.	I am unclear as to what parameters define "dwell time." May want to consider modifications to section 5.2.6.5 as well, to improve clarity. I also am unclear as to how the PCF environment affects fragmented MSDUs.
5.5	Renfro	T		Fragmentation should only apply to either directed messages or broadcast/multicast messages with To DS bit set. For non-ACKed messages, better probability of success will be achieved if the message is not fragmented. The cost of not fragmenting will be that for long broadcast messages a station will not be able to send a portion of the message before a hop and the remainder afterwards.
5.5	Wim Diepstraten	T E T	Delete the last sentence of the first paragraph. The second paragraph below figure 5-25 needs to be made consistent with section 4.1.2.4. It should be specified somewhere that every fragment except the last fragment of a MSDU should have an even Byte length.	Systems other then Frequency Hopping do not have a "dwell time" limitation. The PCF and the Beacon generation is specified such that a normal defer occurs when the medium is busy at that instant of time. This will eliminate fragment concatenation allignment problems in an implementation.
5.5 (global)	Fischer, Mike.	E	change ÒMSDU IDÓ to Òdialog tokenÓ	consistency with chapter 4
5.5, 2nd paragraph	Fischer, Mike.	T	Replace first sentence with ÒThe payload of a fragment shall always consist of an even number of octets except, if necessary, for the last fragment of an MSDU.Ó Also, the middle sentence should state Ò. . . minus MAC header, minus IV and ICV if WEP=1, minus CRC.Ó	simpler implementation, also this provision was approved in a motion at the November, 1994 Plenary Meeting, but the relevant text updates overlooked this paragraph
5.5, paragraph 4 through paragraph 9	Fischer, Mike.	T MAJOR ISSUE	I recommend that this whole discussion of fragment size variation for dwell boundary optimization be eliminated, and replaced with something to the effect that ÒFragmentation shall only be applied when the MPDU required to hold the entire MPDU exceeds aFragment_Threshold. When fragmentation is applied, each fragment shall have a payload length of aFragment_Payload octets, except the final fragment, which may have a shorter payload.Ó	The fundamental reason that fragmentation was added to the MAC was because certain PHYs were unable to deliver maximum length MSDUs in a single PhPDU. This can be overcome using fixed size fragments. The concept of dwell optimization is unnecessarily complex, only beneficial to the FHSS PHY, if at all, and complicates buffer management at the receiving station. The complexity penalizes all MAC implementations whether or not they can attach an FHSS PHY. The benefits are dubious, because if the fragmentation decision must be made based on the amount of time expected to be left after the Ack to the previous fragment, in order to build a MAC header and TXVECTOR for the correct length fragment, but if deferral is needed due to a CCA event, or retransmission of the previous fragment proves necessary, the time calculation is invalid. Finally, with a maximum MPDU size of 400 octets, the FHSS PHY whether operating at 1Mbps or 2Mbps, stands to gain, best case, less than 80Kbps of aggregate raw data transfer, assuming perfect dwell optimization, no extra deferrals, no failures to acknowledge, perfect hop synchronization, etc.
5.5.	M. Rothenberg	T	The Fragmentation Mechanism must be changed to a Window-based, Selective Retransmission Algorithm	1. The current Fragmentation Algorithm is inefficient, adds an overhead of SIFS + ACK time (about 265 microsec in FHSS) for each fragment. 2. The current Fragmentation algorithm is broken: Different aMaxTransmit_MSDU_Lifetime and aMax_Receive_MSDU_Lifetime may cause one side (e.g the receiver) to drop the MSDU but continue acknowledging the following Fragments, hence the frame will be discarded without the transmitter noticing that.

5.6	Bob O'Hara	E	update to reflect new sequence control semantics.	
5.6	Fischer, Mike.	E	change ÒMSDU IDÓ to Òdialog tokenÓ fragement numbers should be 0Ðorigin (0,1,2,3, ...) ÒOnly the last fragment or only fragment of an MSDU shall have this bit set to one.Ó	consistency with chapter 4
5.6	Jim Panian	E	Specify that the duplicate fragment is acknowledged even if the fragment is discarded.	The text does not describe if an ACK is returned for a duplicate fragment.
5.6	Renfro	E	Update MSDU ID to reflect Sequence Control Field. To last sentence add '...but still ACK frame.'	
5.6	Rick White	E	¶ 1: Change MSDU ID to Sequence Number.	MSDU ID no longer used.
5.6	Rick White	E	¶ 4: Change MSDU ID to Sequence Number.	MSDU ID no longer used.
5.6	Tom T.	E	Change 'MSDU ID' to: 'Dialog Token' throughout this section.	
5.6	Wim Diepstraten	E	Section should be updated in its use of the MSDU-ID.	
5.6	John Hayes	E/T	TBD	The current wording describes reassembly as a function of the recieving station. Because it is possible that different APs along the way will have different values for aFragmentation_Threshold that a single fragment will not be able to pass through without additional fragmentation. The current fragmentation scheme does not allow for recursive fragmentation. Therefore, this requires that reassemble be accomplished at each intermediate AP.
5.6	bdobyns	T	An implementation whose PHY MIB parameter aMPDU_Minimum is greater than 2304 plus MAC Header may choose to not implement fragmentation on either transmit or receive.	
5.6	Bob O'Hara	T	Define all atributes in the MIB in section 7	These attributes are not defined.
5.6	C. Heide	t	remove references to MSDU ID	MSDU ID undefined
5.6	C. Thomas Baumgartner	t	Authors of this section need to get with authors of frame format section and decide where the MSDU ID will be.	No MSDU ID in section 4 frame format description
5.6	Geiger	T	Reassembly The description of the contents of a Data Frame header in section 4 are not consistent with the MSDU ID, Fragment number and Last Fragment indicator.	Make both these section agree which ever is the last agreement.
5.6	Tim Phipps	T	Dialog Token: This field allows the destination station to check that all incoming fragments belong to the same MSDU. Fragment Number: Fragments of an MSDU are numbered sequentially, starting at zero.	MSDU ID no longer exists, dialog token is the correct term. All other fields in the specification start at zero, a normal convention in the field of modern computing. It is perverse to have only one field starting from 1. Either it is more sensible to start all fields from 1, or it is more sensible to start all fields from zero. I believe that zero is more usual.
5.7	bdobyns	E	One or more examples of MultiRate frame exchange would be lovely, especially if it showed how the duration field was (correctly) calculated and the NAV maintained. An asymmetric exchange in particular would be amusing.	Not the algorithm for selecting the rate - assume that the STA have already chosen the desired TX and RX rates and that the rates are known to both.
5.7	Mahany	E	Last paragraph, first sentence, replace "on" with "at"	Readability
5.7	Wim Diepstraten	E	Add text that explians how rate capability information is distributed.	
5.7	bdobyns	T	change (two places) "transmitted on STATION_BASIC_RATE" to "transmitted at one of aBSS_Basic_Rate_Set"	aBSS_Basic_Rate_Set is a PHY MIB parameter, while STATION_BASIC_RATE is not.

5.7	bdobyns	T	This section should specify and clarify the use of the defined variables in the PHY MIB section 9.1.1.2 agPhyRate_Grp: aSupported_Rx_Rates, aSupported_Tx_Rates, aBSS_Basic_Rate_Set, aStation_Basic_Rate, aExtended_Rate_Set, aPLCP_Rate, aPreferred_Tx_Rate, aPreferred_Rx_Rate	The IR PHY is asymmetric - it may receive at rates which it cannot transmit on.
5.7	Bob O'Hara	T	Delete this section	Multirate support incurs complexity not commensurate with the theoretical gain in throughput.
5.7	C. Heide	t	remove this section.	(1) there is a great deal of information which STAs are required to interpret in every frame (not just control frames) to make this protocol work. This is broken by multirate support.
5.7	C. Thomas Baumgartner	t	Someone with better understanding of protocol than I should be asked to determine if this section has listed all the frame types that contain data that every other station needs to hear.	What about the End_CF frame? I'm sure that is a frame type not listed here that must be sent at basic rate. There are probably others.

5.7	David Bagby	T	**~~4. Multirate Support~~** Please refer to my comments annotated as "one band = one phy" for background to this comment. The same leadership problem which has resulted in that situation also resulted in the mis-guided desire for multiple rate support. The unpleasant history (as this reviewer understands it) is: The subject of multiple rate support first arose within the DS PHY sub-sub-group. Members from companies participating could not decide whether to support 1mbs or 2mbs for a data rate. Instead of resolving this difference they decided to simply say that they would do both. From a market standpoint this is foolish as the market is conditioned to desire the highest rate possible (all other factors being held constant). In the mean time the members interested in FH PHYs could also not decide on a basic data rate. This resulted in a splintering of the FH gang into two sub-sub-groups which have generally been called the FH group and the hi-speed FH group. Again, the rates involved are 1mbs and 2mbs respectively. This creted a situation where there were people interested in 2 different phys each at 2 different rates all in the same band. While this interest is ok for investigating differences between the proposals, it never should have been encouraged to continue and result in multiple conflicting phy proposals within the draft.	See imbeded comments and annotations

5.7	David Bagby continuation	T	Because each of these small groups is afraid that they don't have enough influence (and hence votes) to get 802.11 to endorse their favored phy/rate combination, they loosely support each other in avoiding trying to create a combined phy proposal for 2.4 Ghz ISM band. The avoidance of the real issue (leadership inability to resolve differences of opinion) has resulted in the side effect of a desire to attempt to support multiple bit rates between stations on a shared media. This is seen my some members as a magic way to enable them to continue doing their own thing. For very good technical reasons no other 802 group has ever attempted to mix different bit rate signals on a single shared medium. In all other 802 groups, different rates are isolated by the physical isolation properties of wired media. This is unfortunately NOT a characteristic of either radio or IR - both of which are inherently shared in nature. The proposals for multi-rate support were consistently rejected by 802.11 until they were eventually watered down sufficiently to appear innocuous enough to finally get a barely passing vote for adoption. The unfortunate results are contained in the D1 draft. It is widely acknowledged (even by the supporters of the proposal) that the performance gain offered by the mixed rate provisions is insignificant to low at best. This reviewer believes that the multi-rate support described in D1 is not only mis-motivated, but also technically unacceptable. The current draft threatens the primary goal of multiple vendor interoperability. The D1 draft says:	

5.7	David Bagby continuation	T	"Unicast Data and/or Management Frames are sent on any available transmit rate. The algorithm for selecting this rate is implementation dependent and is beyond the scope of this standard." Leaving the algorithm for selection of data rate as implementation dependent is un-acceptable as it will result in huge interoperability problems. Therefore, this reviewer will not vote to forward the draft for sponsor ballot until the mis-motivated multi-rate provisions are removed. ~~The following set of rules must be followed by all the stations to ensure coexistence and interoperability on MultiRate Capable PHYs.~~ ~~All Control Frames (RTS, CTS and ACK) are transmitted on the STATION_BASIC_RATE (which as specified before belongs to the ESS_BASIC_RATE) so they will be understood by all the stations in the ESS.~~ ~~All Multicast and Broadcast Frames are transmitted on the STATION_BASIC_RATE, regardless of their type.~~ ~~Unicast Data and/or Management Frames are sent on any available transmit rate. The algorithm for selecting this rate is implementation dependent and is beyond the scope of this standard.~~	

5.7	Dean Kawaguchi	T	Multirate Support The following set of rules must be followed by all the stations to ensure coexistence and interoperability on MultiRate Capable PHYs. ~~All Control Frames (RTS, CTS and ACK) are transmitted on the STATION_BASIC_RATE (which as specified before belongs to the ESS_BASIC_RATE) so they will be understood by all the stations in the ESS.~~ ~~All Multicast and Broadcast Frames are transmitted on the STATION_BASIC_RATE, regardless of their type.~~ ~~Unicast Data and/or Management Frames are sent on any available transmit rate. The algorithm for selecting this rate is implementation dependent and is beyond the scope of this standard.~~ Management Frames are sent at the ESS_BASIC_RATE to enable stations to determine its compatibility and associate or decline association. All other frames are sent at the BSS_RATE. A BSS associated with a particular AP will have a BSS_RATE defined by a management entity. A station attempting to enter the BSS must determine if it is capable of communicating at the BSS_RATE before associating.	Although implementations need not be defined, the standard should include the mechanisms to allow all multi-rate compliant devices to determine when it can switch to higher rates. The current text does not provide any general algorithm nor the mechanisms to enable it to do so. The one dynamic switching method proposed had a patent infringement issue which the committee chose not to tackle. In light of these problems, the only alternative that can be sufficiently defined for a standard is the non-dynamic, management-defined method of one rate per BSS. The text defines the basic method with mechanisms for roaming and CSMA protocol with non-multiple rate units. Note: Both FH and DS PHYs send preamble and PLCP header at the basic rate of 1 Mbps, even on 2 Mbps packets. Thus, all stations are capable of hearing the preamble and PLCP header which contains the length of the packet, i.e., a ÒPHY NAV.Ó
5.7	Fischer, Mike.	T	last paragraph, change Òany available transmit rateÓ to Òany rate available at both the TA and RA stations. If RA capabilities are undetermined, the transmit rate shall be the STATION_BASIC_RATE.Ó	completeness
5.7	Geiger	T	Unicast Data and/or Management Frames are sent on any available transmit rate.	Management Frames must be sendable at the Basic Rate but can optionally be sent at any bit rate. How could you associate with a LAN or set up connections with Basic rate only nodes. I believe that the algorithm used to set the rate can be buried in upper layer management. Unfortunately, I also believe that for purpose of managing the polling list and QoS of the PCF, the bit rate in the CF must be predefined at the time when setting up a connection or the maximum channel usage set at the basic rate and the nodes can optionally send at the higher rate. This must be used by the connection management entity
5.7	Jeff Rackowitz	T	Eliminate this section.	I don't believe that 802.11 should support packets at variable rates in a given BSS. 802.11 radios should be set to a given rate in a particular BSS.
5.7	N. Silberman	T	Re:Multirate Support: Allow support for homogenous high data rate Networks in places where feasible.	Current standard supports only low data rate networks or mixed "speed" networks. In places where high data rate only is feasible, high speed networks will have to slow down the header part lowering the network throughput accordingly. "Mixed Mode" shall be requested only in places where 1 and 2 Mbps stations exist or are expected to communicate.
5.8	A. Bolea	E		Is this section really necessary? It does not add much to the understanding of the protocol. If we decide to keep this section, we should the same format for the MAC and PHY state machines.

5.8	A. Bolea	E	"channel **inoder** to provide some **if**" should be "channel **in order** to provide some **of**"	
5.8	Bob O'Hara	E	delete the first and third sentences	
5.8	Geiger	E	inoder	in order
5.8	Joe Kubler	E	last sentence "inoder to provide some if" should read "inorder to provide some of"	
5.8	Miceli	E	"The MAC management state machines...in order to provide..."	typo in current text
5.8	Renfro	E	Change 'inoder' to ' in order'	
5.8	Tom T.	E	Spelling in last sentence: in order, some of	
5.8	David Bagby	T	Section 5.8 is out of date with cmtee decisions (CFA etc.) and the rest of the document. update for consistency before sponsor ballot.[DB18]	See imbeded comments and annotations
5.8	Geiger	T	General	Where are state diagrams for PCF polling frames and response to PCF polling frames.
5.8 1.1, 2.4.2, 3.2,	Jim Panian	T	Provide MAC service primitives to facilitate the three distribution system services: • Association • Reassociation • Disassociation - including the detection of link outage The above mentioned MAC service primitives will feed into the Association, Reassociation, and Disassocation services in the state machine descriptions as well.	Enough detail must be provided by the 802.11 standard to facilitate hand-off mechanisms on the distribution system.
5.8	Joe Kubler	T	CF FSMs are not in document	
5.8	Renfro	T		Delete state diagram section from standard. It does not add much information and has several inaccuracies. If section is maintained, change format of state diagrams to match those used by PHYs.
5.8	Rick White	T	The MAC state machines are incomplete and must be completed.	There is nothing in the transmit, receive, or control state machines concerning fragmentation. The state machines do not define any of the contention free services. There are references to MPDU ID which no longer exists.
5.8	Rick White	T	Must define the MAC Management State Machine.	MAC Management State Machine is not defined.
5.8	Tim Phipps	T	*Delete this section and all sub-sections.*	This section duplicates information contained in the rest of the specification, at best it is redundant. But, this section has not been maintained (for example: fragmentation has not been included, contention free exchanges have not been included, power saving mechanisms have not been included) and so it contradicts parts of the specification that have been maintained. This section would need a massive amount of effort to be made correct, since this is not realistic I recommend that we delete the section and rely on the main body of the text as the definitive specification.

5.8	Wim Diepstraten	T	This entire section need updating to represent: - Fragmentation functionality - MPDUID handling changes - Add the PCF specifics - Better control Timing aspects - Include Access with prior backoff functionality. - Correct functionality problems - facility to start with backoff if a frame is queued just after a previous transmission by the same station. - Transition C07 should change to C08 in figure 5-29 - add "Backoff=1" to transition C40.	A major update is needed to bring the statemachines up to date with the added functionality.
5.8, also 5.1.5, 5.1.6, 5.1.7, 5.2.10	Fischer, Mike.	T MAJOR ISSUE	The MAC state machines are incomplete and severely out of date with the MAC descriptions. Because this table format does not lend itself to embedded graphics, and the amount of material to be updated, my recommendation is that section 5.8 be replaced with the corresponding material from document 95/14. For consistency, and updating to match MAC changes made since midÐ1994 and appearing elsewhere in this draft, portions of the other listed sections are also updated in 95/14.	The existing MAC state machines are sufficiently out of date as to add negative information content to the draft. They also do not, as currently structured, deal consistently with the use and measurement of time. Accordingly, they are more easily replaced that edited. Document 95/14 is a replacement number for document 94/253 which was not ready in time for the November, 1994 meeting and which was renumbered at Vic HayesÕ request due to the year change.
5.8.1	A. Bolea	E	"transition a listed" should be "transition are listed"	
5.8.1	Renfro	E	Under convention 2), change 'in order to take a transition a ...' to 'in order to take a transition are ...'	
5.8.2	Jim Panian	E	Explicitly show MAC and PHY service primitives driving the flows in the MAC layer state machines.	The MAC layer state machine should be driven by MAC and PHY service primitives.
5.8.2	Bob O'Hara	T	This section must be updated to reflect the current operation of the MAC	Out of date
5.8.2.1	A. Bolea	E	references to PHY end delimiter should be deleted(including figure 5-27). references to MPDUID should be corrected.	
5.8.2.1	Dean Kawaguchi	E	Transmit State Machine All of figure 5-27 is performed in the PHY layer. Refer to para 10.3.3.1.1.	MAC state machine shouldnÕt perform PHY functions.
5.8.2.1 Fig 5-27	Jeff Rackowitz	E	The Tranmit State Machine does not include the current PHY PLCP header definition. i.e. signal field, sevice field, length field, CRC 16. Also, non of the PHYs have PHY-specific trailers so State T6 is unnecessary.	
5.8.2.1	Stuart Kerry	E	Transmit State Machine All of figure 5-27 is performed in the PHY layer. Refer to para 10.3.3.1.1.	MAC state machine shouldnÕt perform PHY functions.
5.8.2.1	Bob O'Hara	T	This section must be updated to reflect the current operation of the MAC	Out of date
5.8.2.1	Geiger	T	Transmit State Machine, Nice try, but the MAC and PHY guys haven't read one another's text.	This state machine doesn't match the state actions the FHSS or DS PHY use. This is too involved to document here.
5.8.2.1	Renfro	T		Prepending (if that really is a word) of the PHY preamble and start delimiter are already part of the the PHY state machines and should be deleted here. Same is true for PHY end delimiter (if anyone ever adds one).
5.8.2.1	Renfro	T		Tx Preamble and Tx Start Delimiter only cover part of the PHY information transmitted prior to the beginning of MAC header transmission. This is all spelled out in the PHY and should not be included here. Also, as defined in the text, State T3 should be Tx MAC header not Tx header since there is also a PHY header. Update reference to MPDUID.

5.8.2.1	Greg Smith	T/E	Between State T3 and T4 there should be the header "PLPC".	
5.8.2.1.	P. Brenner	T	Remove the Transmit State Machine	The Transmit State Machine described here includes mainly PHY functionality
5.8.2.2	A. Bolea	E	references to MPDUID should be corrected.	
5.8.2.2	A. Bolea	E	references to MPDUID should be corrected.	
5.8.2.2	Jeff Rackowitz	E	This state machine needs some work. i.e. RTS and CTS do not include MPDU_ID, R4 does not seem to cover fragmented data, and State R6 should be CF END Received.	
5.8.2.2	Renfro	E		Update all references to MPDU ID. Change To AP to To DS.
5.8.2.2	Tom T.	E	Remove first sentence.	
5.8.2.2	A. Bolea	T		There are missing data types. We no longer check for MPDUID(Sequence Control). Rather we check to see if the destination address is station address. In RTS case, when my_addr=0, all RTS timer references should be deleted.
5.8.2.2	Bob O'Hara	T	This section must be updated to reflect the current operation of the MAC	Out of date
5.8.2.2	Geiger	T	Receive State Machine	This body insisted in several votes to make this state machine and the transmit state machine byte oriented. This state machine indicates the passages of a PDU, not bytes.
5.8.2.2	Greg Smith	T	Fig 5-28 'Receive state machine' R30a and R50a should say: MPDU_ID<>Original_ID. CTS and ACK do not have MPDU_ID fields (see 4.2.1.2)	
5.8.2.2	Renfro	T		Unitdata is not previously defined and not defined in type field. Also, missing from figure 5-28.
5.8.2.2	Renfro	T		Several frame types missing (e.g., beacon, poll, CF Ack,...).
5.8.2.2	Renfro	T		Under **R20a, Other_RTS,** no reason to set RTS timer if an RTS is received. If station is DA, send CTS. If station is not DA, set NAV.
5.8.2.2	Renfro	T		Move 'The CTS timer shall be stopped.' to under **R30b,CTS_Complete.**
5.8.2.2	Renfro	T		Under **R50a,Other_ACK**, the statement that the NAV shall be updated to indicate that the network is now free is wrong if fragmentation is used.
5.8.2.2.	P. Brenner	E	Update the Receive State Machine to reflect the new frame format	
5.8.2.3	Greg Smith	E	Fig 5-29 Bracketed expressions should be used	A&B|C may mean (A&B)|C or it may mean A&(B|C)
5.8.2.3	Greg Smith	E	C00a,b RTS timer and _timeout should be CTS timer and _timeout	
5.8.2.3	Renfro	E		Update references to MPDU ID.
5.8.2.3	Bob O'Hara	T	This section must be updated to reflect the current operation of the MAC	Out of date
5.8.2.3	Joe Kubler	T	retry procedure should allow use of frequency/spacial diversity by allowing interleaving of mpdu to different destinations when retrying.	especially in the case of FH, if a retry can be delayed until after a hop, the MPDU may get through when a momentary channel outage has occured. This allows better utilization of bandwidth as well since communication to other stations (which potentially are not experiencing an outage) can proceed instead of making futile retries to the "bad" station.
5.8.2.3	Joe Kubler	T	fragmentation is not illustrated in the FSM	
5.8.2.3	Renfro	T		Poll, beacons, etc. are missing from state machine.
5.8.2.3	Renfro	T		Delete C00d. If station receives a valid RTS in response to a transmitted CTS, it should transmit Data frame after SIFS time independent of medium. There is not sufficient time in SIFS to perform CCA. Also, C07 should never occur.

5.8.2.3	Renfro	T		In C34, transition should occur when either unicast data was transmitted or any data frame destined for AP.
5.8.2.3	Renfro	T		In C8, Select Backoff, backoff must be integer number of slot times.
5.8.3	Bob O'Hara	E	delete	should now be in section 7
5.8.3	Greg Smith	E	Where is this section	
5.8.3	Jeff Rackowitz	E	Add notes about intentionally left blank or To be specified.	
5.8.3	A. Bolea	T		Text is missing.
5.8.3	Bill Huhn	T	The MAC management state machine needs to be defined.	The current spec leaves this state machine completely undefined.
5.8.3	David Bagby	T	**5. MAC Management State Machines** Missing section - must complete before sponsor ballot	See imbeded comments and annotations
5.8.3	Geiger	T	All management state machines seem to be missing.	It is hard to tell if having no state machine is better or worst than having one which is there but wrong!
5.8.3	Lee Hamilton	T	The Mac Management State Machine must be defined.	Can not implement a MAC without knowing what the management state machine is.
5.8.3	Lewis	T	add MAC management state machines	
5.8.3	Mahany	T	Must be Completed	Omission
5.8.3	Mark Demange	t	Mac management state machine needs to be defined.	Undefined MAC management state machine is inappropriate for a standard.
5.8.3	Renfro	T		Missing. (As all state machines in standard should be.)
5.8.3	**Siep**	**T**	**MAC Management State Machine[must be specified]**	**A standard must be complete in order to be functional.**
5.8.3, 6.2, 6.3, 6.4	Paul Pirillo	E	Define these items, integrate the information into other sections of the document, or delete these sections.	I don't feel comfortable recommending the draft for approval with these sections blank. I am marking this as an editorial comment since I don't have specific corrections to offer.
5.8.3, 6.2, 6.3, 6.4	Paul Pirillo	E	Define these items, integrate the information into other sections of the document, or delete these sections.	I don't feel comfortable recommending the draft for approval with these sections blank. I am marking this as an editorial comment since I don't have specific corrections to offer.
6.	C. Thomas Baumgartner	t	Need to write this chapter so that interoperability is assured and compliance can be determined	Can an 802.11 implimentation be interoperable without this chapter?