

Tentative MAC Minutes Tuesday, March 8, 1994

The meeting was called to order by chairman Dave Bagby at 8:30 AM. Carolyn Heide secretary.

Announcements

- (1) Please register and pay dues.
- (2) When signing book as new participant make sure to add email address.

General

Agenda: last meeting we said we would try to aim papers at about 1 hour each. We have about 14 papers.

Major goals: Settle new functionality set; MAC/PHY interface complete; Simulation group report; MAC standard editing report; due diligence - closed issues versus foundation; detailed schedule for MAC.

Plan:

Monday

- MAC tutorial (done)

Tuesday

- agenda approval
- minutes approval
- simulation group report
- MAC standard editing group report
- new functionality determination

Wednesday

- MAC/PHY interface

Thursday

- due diligence - closed issues versus foundation
- detailed schedule for MAC
- MAC subgroup report review

Rick White points out that Thursday is full working group all day. So we will just work in this order and make as much progress as fast as we can.

Scheduling of papers is first ones of which the chair was notified of in advance, then ones brought to the meeting. The chairman has asked many times for people to make sure they notify him in advance of papers and everyone did very well this time.

Order of papers by subject:

- (1) those that identified problems in foundation (none)
- (2) those that propose new functionality (94/76, 94/37, 94/77, 94/44, 94/38, 94/22, 94/39, 94/40, 94/59, 94/58, 94/60)
- (3) Improvements to the foundation (94/60, 94/41, 94/43)
- (4) Others (94/46, 94/47). Chairman announces Chandos Rypinski's papers are available, but he regrets he cannot be here. He wanted to give a talk in the evening, but wants everyone to know that he would be available to talk with anyone between meetings and will be at the next meeting.

MAC Protocol Version Control, P802.11-94/76, by Dave Bagby

Although the paper doesn't specify it, believes 4 bits will be enough for version.

Questions:

Kerry Lynn: why is version number in frames necessary in wireless when it hasn't been in wired standards?

Dave B: don't believe we will get it perfect the first time. Revisions will be desired. Will need some way to do it reasonably, so let's start from day one. Small investment now to save a lot later.

Mike F: agrees, has used this himself, four bits has been enough even after 10 years of revisions. The other four bits - might be useful to define a range of supported versions - lowest and highest version that a station supports. Not vital, but if you have half an octet free, could be useful.

Dave B: attempted to leave this open for modification later without being a 'new function'.

Kerry: how will we proceed with the draft from here - keep MAC relevant sections in here and vote into draft later, or vote like this a paragraph at a time. If we vote for this are we voting this into the working document, which later gets approved by the group as a whole?

Dave B: I think - purpose is to talk about new functions. Here is a piece we need. Could of just put a motion "let's add version control". But thought explicitly telling the editors what to add was the right thing to do. If MAC group decides we want something it still needs to be approved at plenary. We shouldn't be worried about vast ramifications of every little change.

Kerry: concerned about what is the most efficient way to get work done by plenary. It takes 75% to undo anything we vote in. Might be more efficient to keep all MAC related sections in this group until we are fully satisfied.

Dave B: any decisions made here we take to the parent body. Don't understand what it is Kerry is worried about.

Kerry: putting things in that might impact other areas when it is difficult to get them out again. Do we keep the thing in the subgroup or not is the concern.

Motion #1: **that 802.11 add the ability to handle protocol revisions by incorporating the following changes into draft standard 93/20/b0:**

- 1) On page 4-1 add the line:
2. A fixed-length Protocol Version field which is invariant across all versions of the standard.
- 2) Renumber the previous line items 2-4 to 3-5.
- 3) On page 4-2 add the Protocol Version field to the header diagram. The field is to be added before (to the left of) the current Fixed Header field and after the P-adaptation field.
- 4) Add the following definition of the Protocol Version field to the text on page 4-2:

Protocol Version

This 4 bit field is invariant in size and placement across all revisions of the of the 802.11 standard. The values of the field are defined as follows:

1-15 Revision level of the 802.11 MAC protocol. The values are assigned sequentially starting with the value 1. The currently assigned values are:

0 Reserved for future expansion

This value will be assigned only after all other available protocol revision values have been used. When the 0 value is assigned, the standard shall also specify an extension of the header to include an additional Protocol version field.

1 MAC protocol as defined by the 802.11 standard version x.xx dated xx-xx-xx <editors to replace the x's with the appropriate information to indicate the first approved standard version.

5) Any other editorial changed necessitated by the adoption of this motion (as determined by the 802.11 editors).

Moved by: Dave Bagby
Seconded by: Leon Scaldeferri

Motion 1 Discussion:

Dave B: prefer not to word smith this, please keep comments technical if possible.

Chris Zegelin: in favor of the concept. Would like to see someone who has time to think about adding an element field that talks about capability. The element field is very open in the way it is defined. It would be easy for a lower version to pass an element without know what it means.

Bob Crowder: has experience with this type of thing. Believes not needed in every frame - perhaps on beacons and association frames only.

Leon: supports need for some kind configuration field. This and a couple of other bits are being sent all the time which only need to be done on association - agrees with Bob Crowder. Same with encryption and compression information.

Motion #2: To amend motion 1 by replacing 'header diagram' with 'beacons probes and association frames'.

Moved by: Bob Crowder
Seconded by: Mike Fisher

Motion 2 Discussion:

unidentified: likes the element field idea. But needs to be on every frame so the other side doesn't have to keep track of version for everyone it has to be in contact with. Also, on every frame would allow a version 4 station to revert back to version 1 to have a particular discussion with a version 1 station.

Dave Roberts: with this amendment we have a very state-full design rather than state-less. If H/W is doing this on the fly it can be done as each frame comes in, which makes life easier for S/W implementation. Requires very few bits as specified. Would welcome a detailed analysis of the beacon/probe idea, but is against the amendment.

Greg Ennis: shares concern about keeping overhead down. As to putting it into specific frames - in adhoc we need to think carefully about which frames. Likes Dave's proposal about having it right up front, allows potential difference among versions on information in the frame layout. Against amendment - needs study as amended.

Dave B: better to not adopt amendment. The reason it was put first in the frame was for parsing - if first and invariant that allows you to change anything that comes after. If you start with the assumption of connection then only in some frames makes sense, but in a shared medium a station hears things from a lot of people - would like to be able to do this on the fly rather than have to keep track for everyone I talk to. If later study comes up with a more efficient way we can entertain that. Haven't thought through which frames if only some.

Bob O: in interest of getting the agenda done we should adopt new functionality and not try to improve it on the fly. Urges that we defeat the amendment. Calls the question, seconded by Chris? (many ayes, one nay, no-one feels need for counted vote)

Approved: 2 ayes Opposed: many nays Abstain: - **Motion #2 fails**

Procedural point about voice votes - if you object speak up from now on.

Motion 1 Discussion (con't):

Pablo Brenner: version is a good thing. Also, you should categorize frames as data and management. Don't think this should be in data frames.

Dave R: calls the question, seconded by Mike Rotherberg (many ayes, one nay)

Approved: 28

Opposed: 3

Abstain: 10

*Motion #1 passes***Fragmentation presentations****Fragmentation/Reassembly at the MAC layer, P802.11-94/37, by Mark Demange****Questions:**

Jon Rosdahl: on slide 15, why ack between 2 and 3, not between 1 and 2?

Mark: this shows a window size of 2. With the current IFS system you can do that without giving up the channel. Not trying to suggest a window size just show how it works with this system. The fragments have separate preambles and headers but are the same MSDU, and the fragment frames have no gap between them.

Sarosh Vesuna: no gaps between the bits could have problems for the PHY. How does this fragmentation work with RTS/CTS?

Mark: RTS/CTS goes before all the fragments.

Kerry: this is different from Rick's (i.e. also Motorola) earlier presentation about fragmentation ... it had SIFs between fragments. Need SIFs to indicate division between frames.

Mike Fischer: your justification for fragmentation was potentially losing information and having to re-transmit - doesn't that argue that the ack should not be sent in the same dwell?

Mark: depends what caused the loss - could have been hidden station not noise.

Sarosh: different fragments of the same MSDU must be collected by the AP and re-fragmented on the way down - requires AP to reassemble all fragments of a complete frame. An AP receiving frames from a lot of stations also going across dwell intervals has to keep state information on each reassembly in progress.

Mark: the relay mechanism by the AP is defined in a separate submission. However, the AP isn't maintaining information on multiple packet reassemblies because the use of IFS gaps will stop interleaving from occurring. Any destination station must reassemble, AP or not.

Phil Belanger: adds one or more octet to ack. Do you have any idea of the dimensions here?

Mark: perception is one octet, 8 fragments is enough for an Ethernet backbone. If token ring, packets used are typically Ethernet sort of size on the backbone. But this mechanism supports larger fragments, but doesn't suggest what the correct number is. Simulation should determine what this is.

Jon R: you are worried about microwave oven interference breaking up a large frame. Wouldn't a gap between fragment frames help there? If there is no gap, why bother breaking it up and adding the overhead.

Mark: that comes into the issue of windows size - the worse the channel the smaller the window size.

Fragmentation, P802.11-94/77, by Dave Roberts

Agrees with Marks's reasons why small frame size is needed. But is fragmentation the way to achieve that?

Questions:

Greg Ennis: SNAP proposal - could it be done at the point of entry into the wireless LAN or does it have to be in every station of the wired LAN that wants to talk to a wireless station?

Dave R: either. But probably in the last step before entering the wireless segment.

Greg: must not require wired stations to take special action to talk to wireless.

Dave R: not arguing for a SNAP-like solution, just for not fragmenting within 802.11.

Bob Crowder: what if we are stuck with 128 or 256 byte frame size due to PHY problems?

Dave R: if the frame size is that low, we might have to re-visit fragmentation - but not until then.

Ed Geiger: many stacks on the market don't support small frames. A fixed frame size in the PHY does restrict the dynamic nature of the PHY. When the microwave oven is on a small frame size would be good, but when off a large frame may be just fine. This argues for dynamic frame size. Microsoft and Novell have been forgiving, other stacks won't support new frame sizes. If you accept Ethernet frame sizes you will have a much greater size.

Dave R: other things, like roaming, about the 802.11 will break those stacks. Using the SNAP-like solution may be the preferable way to deal with those. A rational frame size might be 600 bytes.

Jon Rosdahl: this is a new frame type. Protocols written to be media specific will have a problem, media independent protocols will not. For example IPX is media independent, it checks the driver for frame type and handles accordingly. (It is in reality the ODI specification that does this in the driver, not really the protocol - the driver specification.) A new media will break media specific implementations.

Greg: if we rely on higher levels to fragments, fragments that appear on the wireless side have a complete network header and that introduces overhead on the wireless side. Don't think we should make mistakes industry has made before. If we can come up with a better approach we should.

Dave R: we should fix the things that are broken. Mechanism are in place to deal with different frame sizes already (for dealing with large inter-networks for instance).

Dave Bagby: 802.11 is a new standard. Consider whether you would be upset if a Token Ring station sent a frame that claimed to be Ethernet. 802.11 should not send a frame claiming to be anything else - it is a new type. There is some potential benefit for doing fragmentation in layer 2, there is some cost too. We need to decide what is the cost/benefit tradeoff.

Pablo Brenner: (1) saying not going to allow bridging to 802.11 networks could force an entire Ethernet network to reduce frame size to our small size, unless you install something like a router.

Dave R: same thing happens when someone plugs an Ethernet into a Token Ring.

Pablo: (2) SNAP solution - who will be the group that will be doing that if we don't?

Dave R: not arguing for that. Saying that something like that might have to be done to support NETBIOS.

Rick White: IPX has a minimum maximum number of frames.

Jon R: there is an historic guaranteed minimum. Let's not confuse bridges and router - Novell has called a drouter a bridge. A router must have a buffer to receive a packet. The first router (called a bridge) had a buffer size sufficient for a 576 byte packet. No matter how many hops, I could get to the end node without dropping a packet if smaller than that. If larger it could get dropped. Bridging is take something unchanged and pass it through - not one bit changes. If you bridge from Ethernet to Token Ring you must change address - you must route. If you go through size changes you must use routers. Some people call them smart bridges - brouter get broken by new frame types.

Rick: sounds like 600 bytes might be the minimum maximum that some of these can handle. That may not be small enough for some wireless PHYs.

Bob C: problem with concept that router solves all problems. Unless we plan to put a router in all stations we must do fragmentation. There is a serious interoperability problem if one set of 802.11 nodes does fragmentation and others don't. If PHY link can do only 128 byte frames than fragmentation must be done.

Dave R: if frame size must be that low, I agree - must reconsider if this is the case. Motorola's paper seemed to think around 512 byte frames would be sufficient.

Bob C: foundation may be broken if it doesn't provide for fragmentation.

Steve Chen: if maximum frame size of 512 bytes or above is acceptable to PHY, we shouldn't do fragmentation. Without that information we are discussing this too early. Bridging is about frames from similar LANs. 802.1b has MAC bridges for Ethernet, Token Ring and FDDI specifically. If you bridge to a smaller frame size LAN, the bridge will drop that frame. Shouldn't impose that on 802.11 mac.

Dave B: say approximately 1/2 kbyte is that an appropriate frame size for 802.11 to adopt. If yes, some argue for let's not pay the complexity cost of fragmentation. If not, if smaller needed, then it is clear that fragmentation is needed. Sounds like to make this decision here about new

functionality we must look at BER rates in the PAR, and determine what is a reasonable frame size. Can't conclude without knowing that.

Sarosh: MAC level fragmentation is a risky business. Must determine where the problem lies. We could specify what we accept and let another layer deal with it. Let existing products handle it - it's part of the DS and we have agreed not to specify what the DS does. This is one of the requirements we have of the DS.

Jeanine Valadez: every AP is a router?

Dave R: no. A DS may have layer 3 entities within it, or it may be implemented from layer entities. That would be a bridged network and that is dangerous.

Jeanine Valadez: an AP is going to join 2 dissimilar networks, it is a routing function.

Mark: the Motorola presentation did not indicate that 512 byte frames would be acceptable - what the optimum frame size is we are not saying. This proposal does address utilizing the bandwidth optimally, fragmentation can be used to do that. As to BER - significant portion of issue is interoperability with other devices like microwave ovens. It's the character of the interference that is important, not just the BER.

Greg: the SNAP idea is intriguing, you should flesh that out some day in later proposal. Is it possible to do that in a way where source and destination SAPs would pass data transparently, still have end to end preservation?

Dave R: could even use the SNAP/SAP pair and specify a fragment protocol. Embed the protocol id - many ways to do it. The idea is to use an existing mechanism.

Fragmentation or Small MTU?, P802.11-94/44, by Pablo Brenner

Paper had not been circulated at time of presentation, but Pablo explained it very quickly and efficiently.

Questions: none

Fragmentation Discussion

Fundamental question is does the group want 802.11 to do fragmentation. We must conclude this. After lunch the chairman asked for any points that have not been made on the subject of fragmentation.

Discussion:

Chris Zegelin: Hasn't seen anything conclusive as to what way he should vote. Wants something easy for the user to install. Needs more study.

Kerry Lynn: In favor of MAC fragmentation because it lends itself to adaptive forward error correction. A node could send additional fragments which consist of additional forward error bits to compensate for poor channel.

Bob O'Hara: Against MAC fragmentation on religious grounds; not ISO compliant. However, fragmentation should be considered in case it is necessary to violate religious grounds. Microwave oven still a problem with fragmentation; there is a probability function that the microwave will interfere with even small packets. We are defining a wireless LAN not an extension of Ethernet; if the existing stacks are broken then they need to be fixed. The sleeping stations can be accommodated in ways other than fragmentation; a simple solution when coming out of sleep for transmit is to listen to medium for longer than max size frame. A claimed benefit is shorter dwell times of FH which gets you away from a bad frequency sooner; gain no net bandwidth because someone else will hop into this band and have same problem. Every station maintains own timing so not really a problem of being on wrong hop. The reasons for fragmentation put forward were weak. If it can be shown that the frame size will be very small or some other reason then we should consider fragmentation.

Barry Dobyns: in favor of fragmentation. Has a single channel PHY and has an implementation of fragmentation MAC that we could make some measurements on if that would help. Does

fragmentation because thinks that all situations are going to be noisy enough that we need some recovery mechanism. Can tune the segment size. Expectation is that it's going to significantly improve our performance.

Mike Fischer: believes that the potential benefits of fragmentation have been pointed out, but the need hasn't been established. As far back as November the element mechanism was discussed as a means which could be used for fragmentation. We should not include something for which there isn't a proven need. Cautions that there are circumstances, error environments, where fragmentation and selective retransmission after ack can cause a net decrease in throughput rather than an increase. Conjectures that the microwave oven interference using a strategy where ack is sent on same hop as data may be one.

Mark Demange: waking up to transmit is a problem - fragmentation is not offered as a solution to that, simply as "it would hurt you less if you didn't use it". Short dwells - not going to get rid of the problem. Benefit is you move out of the interference more quickly, and delay for that particular transaction goes down. It is a short term delay issue, not a utilization of current bandwidth problem. Waking up on wrong hop - did not intend to state that the system becomes un-synchronized - this is an issue that is very short term. Start times could be off by a worst the length of a stretched dwell. Once cleared everyone is synchronization.

Dave Bagby: torn about the subject. See benefits and costs, so the question is are benefits worth cost. Is an AP a bridge or a router has been asked - this is an implementation minded question, not good for standards development. Some people say there is a bunch of stuff in the field that won't work right - others say you're building something new, so that will happen. It depends what business you're in. Personally thinks layer 2 things should be done in layer 2, similar to discussion on encryption and compression. The cost of this seems to be a software versus hardware point of view. Buffer management seems easy from a S/W point of view, but significant from a H/W one. If it turns out that we can get a reasonable reliability with a frame size that comes close to major protocols that exist, we should count our blessings and do that. But if it comes out to smaller we may have to do fragmentation. We should say we would like to avoid doing this but if the PHY layer needs us to do this we will have to. If it turns out that we have to fragment later we would be solving a problem, not adding new functionality.

Pablo Brenner: similar view. We are talking about PHY capacity, network requirements and cost. If by November we realize we need to add this we can - trial and error approach. Do simulation as soon as possible and/or ask the PHY guys for advice.

Dave Roberts: to Kerry about FEC and that sort of thing - heard we need to accept an Ethernet size frame. Is not adverse to FEC, but is averse to providing fragmentation just to appease an upper layer because of what it thinks our frame size should be. We should not be deciding our frame size based on the propagation speed on Ethernet. If we find a reason at the PHY to do this, ok. Your (Apple's) scheme is quite usable for that, simpler than Motorola's, which can lead to out of sequence packets and all that.

Sarosh Vesuna: do we need fragmentation to make the whole system work - yes. Do we need it at the MAC level - no. We need to deliver packets from station to AP. We are not trying to define a whole system that works - that is up to each vender. The standard for an interface that works is our job, vendor products facilitate the interface.

Greg Ennis: easy connection between LANs has been a goal for IEEE. What happened with different bit order of Token Ring was very unfortunate, has given interconnect headaches ever since. We should try to avoid that mistake. We should make every effort to make what we designing as inter-connectable as possible. We the wireless industry are going to have to do something special to deal with this fragmentation issue. We should not have to change to software that exists to deal with wireless. This is analogous to the security issue - we have exposures that need to be solved but all agree that needs to be done without changes to the wired LAN. Solution needs to allow interoperability without wired LANs have to use a special packet size. Also thinks it needs to be independent of higher layers - there are multiple layers up there we need to support and we need to solve this once and for all for everyone. This is required for the good of the wireless industry, but torn about whether it belongs in the 802.11 or in the 2.5 layer sort of thing.

Might be worth focusing on the difference between the two approaches. For instance, the 2.5 layer suggestion doesn't allow optimizing dwell time while the fragmentation approach does. Not ready to decide.

Mark: MSDUs cannot get out of sequence because of the Motorola proposal. Feels that it simplifies things to allow things to go across the channel out of order, but if this is too bad, it can be modified easily.

Phil Belanger: complexity is the issue. Whether or not to fragment is based on what the PHYs can support. What is unique about these PHYs that causes us to need short frames and how short are they? Maybe not in the MAC, but it has to be solved. Could be swayed either way right now. What the PHYs can support is how the decision should be made. The complexity of Mark's system was overblown by some people. It is quite simple, in general only one MSDU has outstanding fragments.

Jon Rosdahl: with AP providing routing type functions, group seems to have decided that AP has some type of wired thing and at the dispersion mechanism of the antenna (DS,FH,IR) the size frame has to be very small. This moves to the conclusion that we must fragment at MAC layer. We are defining MAC to MAC communications here, not MAC to outside our world. The mobile client is not going to have a lot of resources to buffer things. The MAC layer's job is to take data and get it to the other guy. We have to compress, encrypt and now we say fragment too. Now the PHY says we'll fragment too, and we wind up with a little stack in every layer. Doesn't think the MAC layer is the place for fragmentation, will need to be done but not there.

Kerry: maximum supported frame size is going to be PHY dependent. If we don't do it in the MAC we will have to export the PHY type above the MAC, or take the least common denominator.

Bob Crowder: issue is what can the PHY support at a reasonable error rate. We have to ask the PHY group. Believe the distribution system is misunderstood - there is no issue of DS relative to frame size. The end wireless node has to deal with it. Clearly it is in the protocol this group defines. Motorola's (i.e. Mark Demange) ideas looks very efficient. If MAC deals with compression and fragmentation, we will put them in some order than makes sense. The MAC is the place to do this since they interact. Not strongly in favor of doing it in the MAC, but the DS can't do it, it is a MAC issue.

Straw poll, who is:

1. confused and would abstain if we voted (some)
2. in favor of fragmentation in the MAC: (some)
3. violently opposed to fragmentation in the MAC: (1)
4. in favor of seeking advice from the PHY group: (a lot)

The group appears not to be prepared to make a decision now. We should specifically direct Kerry Lynn to bring this up at the MAC/PHY interface meeting tonight.

But Dave is loath to say decide later because we are on a tight time table. Let's move this to things we should do later in the week. No one strongly objects to this.

Laura Hindy: thinks it's took complex to deal with in a 3 day period.

Mark: view that it's not just a PHY issue here. They may say we can send 12 msec duration frame - we need to operate in a world where other unlike systems are there.

Frame Windowing at the MAC layer, P802.11-94/38, by Rick White

Questions:

Dave Roberts: affect on NAV?

Rick: needs to be worked through in more detail.

Dave Bagby: for both the fragmentation and windowing submissions - what do both/either allow me to do that I couldn't before?

Rick: these improve performance/throughput.

Dave B: would you propose windowing without fragmentation?

Rick: I wouldn't rule it out, necessarily, it is independent of fragmentation.

Dave B: will put a decision on windowing right after the fragmentation decision tomorrow [no one objects].

Bob O'Hara: time gap between fragments (none shown no drawing)?

Rick: yes, shorter the SIFS.

Laura Hindy: don't see any quantitative analysis to show this performance enhancement is necessary.

Rick: correct - simulation should be done.

Simon Black: would like to see a benefit analysis on this proposal. MAC is already complex. Cost/benefit analysis is very important when just a performance improvement is suggested. Intuitively feels the benefit is small versus the complexity cost.

Sarosh Vesuna: windowing has benefits for small frame traffic between two nodes. Multiple frames piggy-backed without having to do CSMA for each.

Mike Fischer: keep in mind when considering windowing/fragmentation issues - if we adopt fragmentation due to PHY frame length constraints and results include windowing, we may have to add MAC/PHY parameters regarding how long a PHY can go without sensing the medium. We might create a situation where multiple small frames could run longer than they should go without sensing the medium.

Wim Diepstraten: header contains indication that ack is required. What if that frame gets lost?

Rick: could go in all the frames. But otherwise it is a loss like any other, backoff performed.

Wim: performance improvement achieved - is windowing required to make fragmentation beneficial? Or is fragmentation without windowing generating overhead that reduces performance?

Rick: if I had to contend for the channel every time, it would add overhead. Windowing with fragmentation is more beneficial than fragmentation without windowing. Window size should be programmable - set to one and you have no windowing. But loss of channel contention is a performance benefit.

Reports missed on agenda earlier

Simulation Adhoc Group Report, by Kerry Lynn

Due to inability to get source code to adhoc group, limited work done. Kerry Lynn can pass out source to anyone who signs a license agreement, tomorrow after the meeting. License agreement basically says Apple is not trying to make or lose money here, and it is not liable for decisions made based on simulation results. Also, you can't take the source code and build it into something and sell it. Any changes go back to Apple so they can be shared with this group.

New functionality added into source: parameterized RTS/CTS; traffic model reflects AP traffic, more client/server oriented; added antenna diversity, Raleigh fading; made it possible to eliminate channel model dynamics, so channel can be kept constant over time.

Discussion:

Jim Schuessler: plans for future?

Kerry: group needs to get together again. Improved backoff needs to be added. Also multi-cell model, nodes acting as interferers.

Chris Zegelin: default parameters need cleaning up, they are out of date.

Editing Group Report, by Jim Schuessler

Greg, Bob and Jim met every week between last 2 meetings. Sections were emailed out on the reflector.

- Worked on sections 4 5 and 6. - would like feedback on editor notes in there.

- Section 6 - MAC management. Used to have a section 10 'layer management' but the document 93/190 functionally didn't seem to fit in section 10. Section 6 is mainly synchronization and power management.
- Identified holes throughout. Editors don't have charter fix these. Such as the subject of addressing, seems to be spread out and needs to be isolated.

Greg Ennis adds that editors notes identify places where there are holes. When people see them they could address contributions to them. These notes are pointers to small submissions.

Discussion:

Dave B: what do we do now that we have this thing? Everybody should read it for "did the editors accurately transcribe into the draft the decisions of the group". That is different than "do I like what's in the draft". That's a different kind of submission for another appropriate time. We need to find mistakes as early as possible.

Jim S: included 93/190, but also the contention free acknowledgment submission by Jim Schuessler and the state machine document.

Tom Baumgartner: is that really the only way to have feedback? One type of feedback is spelling, grammar - we could just give those to editors. Everything else goes into a submission? How do we get our red marked comments across if submission the only way?

Dave B: go to editors and ask is this an editorial problem. If they agree they can change it.

Tom B: go to the editors and try to work it out with them, and if that doesn't work bring a submission?

Jim S: expected that there would be a large block of time to go through the standard line by line at each meeting.

Dave B: 150 people trying to agree on wording is not appropriate.

Jim S: that places editors in the position of deciding whether changes are editorial or technical.

Dave B: accept the responsibility you took on.

Greg: written comments identifying section number would help the group. Unfortunately the document is only available in PS so that can't be done electronically.

Dave B: people can get together with the editors and decide how this is going to be done.

P802.11-94/22, by Kerry Lynn

Kerry suggests that people should read the ETSI security document - you can get it from Kerry if you want it.

Paper not circulated at time of presentation. The gist of it was that 802.11 should adopt the encryption algorithm developed by Apple which uses RSA RC4, as the minimum security level required. Kerry described the generally how the algorithm works. RSA agrees to make RC4 available for a reasonable fee if it is accepted as the mandatory minimal security algorithm for 802.11. The fee ranged from 50 cents in single quantity, in stages down to 12.5 cents for large quantities.

Discussion:

Laura Hindy: is there any way we could get a non-proprietary, free algorithm?

Kerry: if you start now you might get an algorithm approved by NSA in 8 to 10 months.

Dave B: the reason that Kerry has done all this, to make an exportable algorithm available - to attempt to do the same thing with a different algorithm from scratch is do-able, but very hard. I don't view this as the world's greatest security, but its pretty good. We want something that gives you the equivalent of the physical security of a wired LAN. Is convinced this is as good or better than the security that comes from a wire.

Laura: we don't want to pay that 50 cents per unit!

Dave B: the guideline previously agreed by the group was keep it under \$100. Don't know of anything that is available for 0 price.

Kerry: the approach here is leveraging off previous work done to make the RC4 algorithm exportable. The president of RSA was going to come to Vancouver but couldn't make it. The price schedule would drop down in high volumes. What is the perceived benefit for the user - can you justify the cost of another .50 cents? I think so.

Dave B: this is the first thing that a real cost has been affixed to.

Jon Rosdahl: does Apple get any of the 50 cents too? RC4 is RSA, but the rest is Apple.

Kerry: that is undecided.

Laura: if we don't want to export why do we care? Why not look for something cheaper, unapproved by NSA?

Jon R: failure to get NSA approval on anything that looks like security may cause your company to lose your export license.

Kerry: if we decide we want an algorithm and we choose one non-NSA cleared we could not include that in hardware that leaves the country. That means 2 versions of everything built, and that is a royal pain.

Greg Ennis: asking us to adopt an encryption algorithm. I'm no expert, but can you illuminate the process Apple went through to choose this algorithm?

Kerry: to speed through the approval process we wanted to leverage off precedence. There is a recipe - use RC4, restrict key size, etc., that expedites process through the system. Started last July, and only now can say it's in the works. RC4 has a good reputation in addition to it being exportable. It is included in some Apple products already, has a name in the industry. The pseudo random number generator is the key to the machine. Wanted to maximize the seed size, but have it bigger than the secret key size. Really relied on the reputation of RC4 which is very good. It's a good compromise.

Laura: do existing LANs generally want to be exportable, and so have this problem?

Dave B: wired LANs do not have the security issue of wireless.

Tom Baumgartner: if you get away from the LAN side and look at things like Lotus Notes - there are people who have products with DES encryption in them and would love to not have this problem.

Dave B: only recently have the export laws been changed. Previously if I had brought some of these programs that include encryption on my computer with me to this meeting I would have committed a crime. They have been recently approved for personal use.

Barry Dobyns: are we considering making encryption an optional part of standard or mandatory? When we write the PIX perform - mandatory or optional?

Dave B: previously we discussed that we wanted all stations to perform a wired equivalent level of security. We have not closed the issue yet because we were waiting for a cost example. The intent is that all stations would support it.

Tom B: was certain that we said that one of the options would always be non-encryption.

Dave B: also a minimum level that everyone must have.

Leon: that issue is still open, and closing this motion would close it.

Ed Turner: is there any place to which exporting this would this a problem?

Kerry: maybe France, and thinks there's a State Dept. license you can get to solve that problem.

Dave B: because France is only country that has import restriction on this.

Leon Scaldeferri: must have an export license to give any non-US citizen the algorithm - they even have a problem about who can sit in meetings when the algorithm discussed. It's not an easy thing to deal with. There are procedures if you create your own algorithm you have abide by. When this was brought up a request was made for candidate algorithms and this is the first.

Simon: implementation size of RC4?

Kerry: inner loop on a 6800 is 96 cycles.

Sarosh Vesuna: (1) how does this apply to product manufactured for sale outside the US?

Dave B: US has most restrictive laws

Kerry: suggesting algorithm for all units so any station can communicate with security if he wants.

Sarosh: (2) what do we get out of this company - algorithm, implementation? How do we make sure it's standard over all copies we buy?

Kerry: object code license prices are what's quoted here. For an addition \$500 or something you get source code. I got the source code and optimized it for my target.

Sarosh: has it implemented in any form of H/W, or has it been evaluated what it would look like to implement in H/W?

Kerry: that is a rubicon that you can cross with NSA. To them there is a difference between a S/W implementation and a H/W one. It changes the device from a LAN adapter to a cryptopographic device. Each of you who are considering hopper implementations should make sure your hop sequences are not user programmable ...

Phil Belanger: you give the cost of a license for software implement - what is cost for H/W implementation?

Kerry: no idea, sure it's just as reasonable. RSA's motive here is more to get related business, more than it is to make a fortune off the RC4 object itself. They don't want cost to be an issue here - if 50 cents is too high I will pass that input back.

Sarosh: impression was that we would allow products that would have different levels of security allowed. Last January we had a lot of security discussion.

Dave B: we adopted 802.10, it maps numbers to algorithms, including a null algorithm. Now we are discussing the minimum tool kit for what has to be in a station.

Sarosh: if this gets into the standard what does it mean?

Kerry: my motion says we have to include this in every node. You can choose not to use it but you have to know how to use it if someone sends it to you. You cannot build a station without this. This is orders of magnitude more secure than a scrambler.

Motion #3: **The RT cryptographic algorithm shall be included as an integral feature of the 802.11 standard.**

Amended by motion #4: **The RT cryptographic algorithm shall be included as an integral optional feature of the 802.11 standard.**

Moved by: Kerry Lynn
Seconded by: Mike Fischer

Motion 3 Discussion:

Dave B: adopting this closes issue 6.10 "Shall the 802.11 standard specify one or more publicly available privacy algorithms which all stations shall be required to support?"

Greg: what form would this take in actual standard text? We would need exactly this algorithm in the standard.

Kerry: 94/22 includes a section toward the that has 'C' code - there is a service interface specification. NSA has a referenced implementation that can be referred to in the standard text.

Greg: is there a precedent in standards for doing this? How do we specify it if we can't look at it?

Leon: there are other algorithms like that in the cellular arena. GSM for instance is a standard, but it is not public.

Simon Black: the document for GSM is standard, you sign up and you receive the document.

Greg: how, mechanically, do I get the standard written if we do this?

Jon R: RC4 is approved by NSA, is the algorithm proposed here as a whole approved?

Kerry: NSA reserves the right to examine every individual product if this is included in it.

Dave B: Kerry has gone through an approval procedure and probably has been told not tell make some things public. Believes he has been told approval would probably be given. But probably has been told that he has to say they can examine everything. That it is probably not their actual intention.

Bob Crowder: if they do test, who pays?

Kerry: it doesn't cost anything.

Jon Rosdahl: level of security this provides - in order to become 802.11 compliant would we get a magic number of security rating?

Kerry: don't know. Believe the system as a whole must be considered. This algorithm provides a security level appropriate for the purpose we are considering.

Frank O'Neill: there is a complexity associated with this, a burden on stations involved. Why is making it 'integral' critical to success? What about an optional scheme, stations negotiate a level of security and level 0 is no security. Stations must have the negotiation ability, but the negotiation can be unsuccessful.

Kerry: I won't drop down to level zero.

Frank: then you can't communicate with a station that doesn't implement this. But two stations that are willing to drop down to zero can. You are saying I can't choose to build a station without this. If I don't want to connect with you, I shouldn't have to do this.

Kerry: security equivalent to a wired LAN is needed for success.

Simon: might be nice to have stations that always operate in the clear. You are saying that you are forced to operate in the clear to talk to him. If optional use, the problem still exists, so why force it.

Kerry: intention is rise to the high water mark.

Simon: if all stations have to implement then no station can operate in the clear because everyone has to have it, so they will always use it. This is implied.

Motion #4: To amend motion 3 to substitute 'optional' for 'integral'

Moved by: Barry Dobyns

Seconded by: Tom Baumgartner

Motion 4 Discussion:

Bob Crowder: there are a lot of applications where this is not a necessity. After that last discussion about approvals need we encryption used, can't help but feel a companies like that Shipstar and Apple might have different levels of trouble to get all approval. Can't conceive of the complex this might make life for us. Can see that we could be precluded from making bus analyzers because of having the algorithm. It is not normal to have make a motion to make something integral - a system won't work without an integral feature. The fact that this motion exists proves it should be optional. Feels it must be optional.

Bob O'Hara: those who don't believe that this feature should be optional are relying on market needs and history. In a mature markets that's a way to extrapolate. We are not in that case - this is an immature market and as such the belief that no security at all, or simple scrambling is sufficient, may be true. Only in rare instances like this meeting might we want a network operating in the clear. Feels that 'optional' is not an option.

Dave B: don't know about supporting this amendment. Honestly believes that an unsecured wireless LAN is probably a failure as a successful product. You must abstain or practice safe data. Would go for a more stringent encryption minimum than this. Would take it as an option if had to, but think that's a mistake. You only get to sell perceived security once. Would like to remove in the clear as a possible operating mode, don't want this as an option. A corporation will not install an unsecured wireless LAN. Don't react emotionally to the price, would take this price over no security.

Tom Baumgartner: don't profess to be able to see forward in time. But commercially right now I'm not experiencing the mindset of the corporation that you are. Won't mention companies we have sold to, but they are in industries where you might think that this is a problem, and without encryption products have been sold. My experience says you are wrong about the present need. Maybe the fact that these are IR, versus RF, is why - a level of physical security is provided with IR. This MAC is supposed to support all PHYs, therefore I support the amendment.

Pablo Brenner: agrees with Tom. We are selling now, a wireless RF bridge and we have an encryption option. The user asks do you have it, we say yes for 10% more, and they say we will let you know if we need it. They don't want it now when they have to pay for it.

Dave Roberts: torn. thinks optional is ok. Don't view this as a MAC interoperability issue. It is a policy decision of an organization to operate in the clear. Operating in the clear is stupid, but should anything else be forced? It can be done in H/W without much cost.

Mike Fischer: agrees that optional is better than none, but failure to make standard encryption method is going to make a mistake. Market now is just for evaluation. People are surprised when they realize it isn't as secure as a wire. People suggesting that user perception is that scrambler is sufficient for that, are hoping that will shield them from product liability litigation.

Bob C: makes bus powered dives (with power concerns similar to battery devices). Typically algorithms repeated are a power drain, no insignificant. This is not a wash, cost is not the only issue. In the industrial world you are normally encased in an area where the energy is inside the building, and a lot of the data is not of the kind that the customer is concerned about. Many networks exist like that now. There is a technical penalty without a clear market reason. Don't do me a lot of favors, if I destroy my implementation that's my problem.

Carolyn Heide: believes that people are being short sighted about the markets for wireless. Corporate LANs, which do require security, are not the only applications. There are applications such as industry and education that don't need security. A classroom with a teacher broadcasting to students, using IR that doesn't go outside the room - why would that need security. What about 'desktop LANs' - why do I want encryption on the IR link between my computer and my pointing device? Also agrees with Tom that one of the PHYs inherently provides physical security equivalent to a wire, it should not be penalized.

Dave R: to Bob C - the power issue is not a big one. The intent of Kerry's motion is good - how do we say that this is the minimum when the minimum is used, when in the clear is not the case. What do we say by making this optional? Have we in effect not blessed this algorithm? Calls the question, seconded by Kerry (ayes, no nays)

Approved: 17

Opposed: 11

Abstain: 7

Motion #4 passes

Motion 3 Discussion (con't):

Dave B: will take this as an option. is a little loath to accept this algorithm without signed letter in front of me with a commitment from RSA.

[sec note: at this point Kerry produced such a letter, which Dave read to the group. I am not a good enough typist to get it as Dave read, but when asked if it was going to appear in the minutes I said yes, assuming I would get the letter from Kerry. On further thought Kerry has decided not to make that letter available for the minutes. Anyone wishing to know the details may contact Kerry Lynn.]

Dave B: read a letter to Kerry Lynn which described basically what Kerry said in his presentation - that RSA would make this algorithm available on a cost schedule which varied from 50 cents in quantity of one down to 12.5 cents in high volume, if this algorithm was implemented as part of the standard.

Phil Belanger: Tom and Pablo have said in their brief marketing experience with wireless LANs says that customers are not requesting security, and if you mention a cost they say "I'll get back on you". The conclusion I draw from this is it ought to be automatic and in there all the time so they don't have a choice or an extra cost. Would have prefer to have this as an integral feature. This motion is bass ackwards. We don't have much information on this. We haven't even seen Kerry's paper. Believe that it unlikely that productive progress is to be made. Call the question, seconded by Tom Baumgartner. (ayes, no nays)

Approved: 8

Opposed: 11

Abstain: 16

Motion #3 fails

Remember the MAC/PHY interface adhoc group meets at 8 PM.

Meeting adjourned: 6 PM

Wednesday, March 9, 1994

Meeting called to order at 8:40 AM, by chairman Dave Babgy. Carolyn Heide secretary.

General

Asked everyone last meeting to put an abstract on first page of all submissions with category - foundation is broken, new function, improvement or other. Based on the abstracts Dave asked the authors of the remaining papers (94/39, 40, 49, 58, 41 and 43) to summarize their papers then asked the group for each if they felt they needed to hear the presentation. The group felt there was a need to here all of them.

Just to try to close the discussion we had on security yesterday, what would the group think of a motion "that we close issue 6.10 with: Yes, and the algorithm for wired equivalency will be RT as described in doc 94/22. The use of this algorithm will be optional and any 802.11 compliant station which implements any security algorithm other than 'none' shall include the RT algorithm in the set of security algorithms supported".

Some people really like. At least 3 really hate it.

Discussion:

Barry Dobyns: objects to RT in particular. Agrees that security is necessary.

Sarosh Vesuna: don't have enough information to make the decision about RT.

Kerry Lynn: not sure the business proposal from RSA applies if the algorithm is optional - economy of scale.

Dave B: ok, won't make this motion, the problem is clearly not as simple as he thought.

Leon Scaldeferri: could people who feel they don't know enough go to RSA and discuss the algorithm with them?

Kerry: yes, he believes so.

**A Packet Delivery and Relay Strategy for the Foundation MAC, P802.11-94/39,
by Mark Demange****Questions:**

Wim Diepstraten: if all stations communicate with each other via the AP, using AP relay at all times, there is never any potential sequencing problem. It is only when some gets relayed and some goes directly that sequencing problems are created.

Mark: yes. There are sequencing issues that this mechanism addresses, there are also sequencing issues in the foundation MAC.

Phil Belanger: what's the difference between the 'to_ap' bit and the 'relay' bit?

Mark: thought the to_ap bit meant the frame was destined for the AP. there might be a simpler way to use the bits, let's continue with the philosophical idea here.

Dave Bagby: you are trying to short circuit the DS function - it will take too long to send the frame into the DS then get it back, so you want to short circuit this.

Dave Roberts: in what case does peer to peer in the foundation MAC have a problem?

Mark: you need a mechanism that allows communication between 2 devices in the same BSS which are hidden from each other. Someone in the same BSA that you couldn't communicate with directly.

Dave R: doesn't the foundation MAC say all frames are relayed through the AP?

Mark: didn't see that so clearly.

Dave R: there is a sequencing problem in the foundation too - someone far away, in a system with a slow DS, is sent a frame which has to go through the DS. While that frame in transit the station moves, and gets something sent right to it. If we are relying on sequencing being consistent we have a problem with roaming.

Mark: maybe that exists, this is not addressing that problem.

Dave R: since this doesn't solve that problem, should we introduce this much complexity?

Motion #5: **Allow direct station-to-station communications in a BSS with an AP.**

Moved by: Mark Demange

Seconded by: Kerry Lynn

Motion 5 Discussion:

Jim Schuessler: this is already in the standard, already supported - chair should this motion rule out of order.

Dave B: he can make any motion he wants. Personally thinks this is not worded as well as it might be, too broad.

Jim S: if we vote against this because it's redundant do we have to remove this function from the standard?

Mark: if everyone feels that this is already in standard I could remove this motion.

Straw poll - one person believes not supported currently, a lot of people believe it is.

Bob O'Hara: calls the question, seconded by Wim Diepstraten (ayes, no nays)

Approved: ayes Opposed: no nays Abstain: - **Motion #5 passes**

Motion #6: **Permit AP to relay a packet ~~only when necessary.~~**
~~- when destination station is out of range of source.~~
~~- when destination station is asleep.~~

Moved by: Mark Demange

Seconded by: Kerry Lynn

Motion 6 Discussion:

Dave Roberts: against, current functionality as specified in the foundation is the appropriate. If the default case is always forward then no one needs this complexity, and stations can make optimal decision when they know what they can do. Requiring all stations to be optimal at all times places unnecessarily burdens some stations.

Leon: issue 11.5 answered yes is what you are suggesting, without the "only when necessary". Supports answer of yes, not sure how I feel about only when necessary.

Dave B: procedural note to issues log editor - in issue log reference to close with yes.

Motion #7: **To amend motion 6 by striking everything "relay a packet".**

Moved by: Sirosh Vesuna

Seconded by: Dave Roberts

Motion 7 Discussion:

Barry Dobyns: calls the question, seconded by Jim Schuessler. (all ayes)

Approved: 17 Opposed: 7 Abstain: 4 **Motion #7 passes**

Motion 6 Discussion (con't):

Dave B: this proposition adds some complexity with some problems, no new functionality. Philosophically opposed to the relay game. This allows communication to station hidden from me, out of my immediate range - that just short circuits distribution systems do.

Bob O'Hara: calls the question, seconded by Barry Dobyns(ayes, no nays)

Approved: ayes Opposed: 2 nays Abstain: - **Motion #6 passes**

Motion #8: Adopt packet delivery and relay strategy proposed in P802.11-94/39 as a basis for inclusion in the draft standard.

Moved by: Mark Demange
Seconded by: Rick White

Motion 8 Discussion:

Motion #9: To postpone motion until right after Pablo's paper.

Moved by: Arthur Coleman
Seconded by: Bob ?

Motion 9 Discussion:

Leon Scaldeferri: calls the question, seconded by Barry Dobyns (ayes, no nays)

Approved: ayes Opposed: 1 nay Abstain: - **Motion #9 passes**

Motion #8 postponed

AP based CTS, P802.11-94/43, by Pablo Brenner

Questions:

Sarosh Vesuna: what about overlap of adhoc and AP?

Pablo: they are not in the same BSS. The adhoc guys don't care about what the AP says.

Dave Bagby: CTS is a clear out the media command from the AP. Suspect that in this overlap case this has a problem.

Pablo: stations working in adhoc when there is an AP - has not been covered by this proposal. Don't believe we have dealt with stations in adhoc network communicating to stations in infrastructure - they can't communicate.

Arthur Coleman: if AP always sends CTS, if the receiving station is out of range of the sender or is power saving, the data frame comes out needlessly.

Pablo: AP picks up frame and keeps them for relay if there is no ACK.

Rick White: requires all data frames use RTS/CTS - standard instead of option.

Pablo: all station to station traffic should use RTS/CTS. Traffic from AP could be without it.

Rick: changes reason for RTS/CTS put into foundation. Was meant to tell the people in the area around each station there is traffic. In this scheme the receiving station didn't send the CTS so the area around it is not cleared out. Stations out of this BSS can now interfere.

Pablo: technically it's not such a big change, philosophically it might be.

Greg Ennis: the advantage of the RTS/CTS in the foundation is it is independent of which BSS you are associated with - it clears the medium for all stations in neighborhood of receiver. What you are proposing looses this functionality.

Bob O'Hara: you have mentioned some things that you say are broken., and some enhancements that are needed. Some intuitively seemed correct, however intuition and networks are often not compatible. Is reluctant to accept this without some simulation results.

Pablo: agrees. Sent email to Kerry about getting this simulated. But if we have wait for simulation results, then we come to next meeting and get told that it's new functionality and it's too late to add it.

Motion #10: To perform modified simulations on the MAC protocol, and if it proves better, then include the changes mentioned in P802.11-94/43:

in infrastructure mode:

- stations do not send CTS to each other
- stations do not update NAV based on other stations messages
- the AP sends CTS for any valid RTS (if the medium is free)

Moved by: Pablo Brenner

Seconded by: Barry Dobyns

Kerry Lynn: friendly amendment to say Pablo will change the simulation himself.

Pablo: rejects amendment, says he can't do it himself.

Motion 10 Discussion:

Dave Roberts: this brings back in a PCF. Makes the AP a PC for the BSS. Blows away the point of RTS/CTS clearing around receiver. Second, does this task the adhoc simulation group to do this? Pablo is welcome to do it himself, but given our limited time they should simulate the MAC we already have. Believes this potential improvement is broken anyway.

Pablo: would do the simulation. Would have preferred a third party as being more appropriate, but will do.

Dave R: why is Pablo making a motion to allow himself to simulate his own modification. Asks him to withdraw motion.

Dave Bagby: personal opinion - thinks there are an innumerable number of problems with this. This changes one of the most fundamental aspects of the foundation. It is a misunderstanding of what RTS/CTS are used for. It assumes adhoc networks don't exist in same space as infrastructure networks. It is not an improvement. It is such a major change I don't know how to speak strongly enough against.

Mike Fischer: not as fundamentally opposed as Dave, but wanted to point out it does make a fundamental change in MAC. That space over which RTS/CTS manages is the infrastructure BSS not the space around the stations. It opens closed issues.

Dave B: if passed, procedural problem. It is not reasonable to reverse sets of closed issues as a side effect of a motion.

Mike F: if we had analytical or simulation evidence that something about the foundation was broken we would have a better reason to do this.

Sarosh Vesuna: performance degradation because two conversations, which don't interfere with each other, but are both visible to the AP, could have gone one simultaneously. With this change they can't.

Point of order, Kerry: motion is to simulate or not, not is the proposal any good.

Dave B: no, the motion has obligations according to the simulation results.

Pablo: doesn't believe in real world applications any degradation occurs.

Sarosh: ACK not seen by AP, and AP ACKs. Says AP ACKs packets where from/to doesn't refer to him at all. The ack may have been sent, but the AP didn't hear, and sends a duplicate.

Pablo: once AP has ACKed how to handle it is now his business.

Sarosh: calls the question, Kerry Lynn seconds (ayes, one nay)

Approved: 2 ayes Opposed: nays Abstain: - ***Motion #10 fails***

Dave request, but no one wants a count on that last vote.

Motion #11: **That we simulate the unfairness problem shown in this P802.11-94/43 and if it proves broken change DFWMAC to solve problem.**

Moved by: Pablo Brenner

Seconded by:

Dave B. suggests that making a motion to do a simulation accomplishes nothing. If anyone is interest they will simulate. As to the second part, we have asked for papers on "if broken show how to fix" - affect of motion is to modify procedure.

Motion #11 withdrawn

General

Phil: clear from last 2 papers there is an area in 93/190 that is not clearly specified - relaying, or how the AP gets involved in the interaction, deciding when to CTS or ACK. Perhaps that's an area that a small group of people could make progress with maybe using some of the ideas from these 2 papers. I believed some people misunderstood things in the paper, then I went to paper and found they were not explained there.

Return to postponed motion 8 discussion

Dave Roberts: ask mover and seconder to withdraw motion in light of Phil's last comments. Motion will probably not pass because of that and this would save us a lot of time.

Mark: ok.

Rick: wants to ensure that this does not preclude this type of new functionality moving forward (in light of the fact that this is supposed to be the last meeting at which new functions can be proposed). A small group to address this is a good idea.

Dave B. believe there is no new functionality here, this provides improvement (in the eyes of the author). No one in group disagrees with that.

Rick: are there people willing to work on this. [At least 8 people say yes] The seconder agrees to withdraw the motion.

Motion #8 withdrawn

Improved Frame Format for the Foundation MAC, P802.11-94/40, by Rick White

Idea of negative ack is not in the current protocol. There are benefits to it which improve performance, and it differentiates between destination station is within range and heard frame with an error, and there was a collision or destination station is not in range. Absence of ack is all we have now, so we cannot make this distinction. In order to have this frame modifications are required.

Questions:

Dave Roberts: do you mean length (bytes) or duration (time)?

Rick: really feels length, but duration was used in the original document, so used that.

Dave R: in the original document we purposely used duration, sometimes the time is important.

Wim Diepstraten: intention is length is length of that particular frame. The duration in the RTS and CTS case means the entire transaction. The RTS/CTS would need both duration and length in what's proposed here.

Rick: yes. Duration should be changed to length in the header.

Greg Ennis: if length is to mark the end of the frame, the PHY groups are grappling with this issue. One hopper proposal says that length information although generated in the MAC and passed to the PHY goes into the PHY frame header. In the case of other PHYs there may be other ways to determined end of frame without length, so maybe it doesn't belong in any frame.

Rick: aware of that, and is undetermined.

Arthur Coleman: isn't that a PHY question?

Rick: not worked out yet.

Mike Fischer: said last night, only FH PHY has determined that it needs the length. DS did say that they could take advantage of it, if there is going to be a PHY CRC protected header. Given that that information is passed from the MAC it could be used.

Simon Black: offering something that is only a performance improvement. Extra functionality without know what the real benefits are without simulation data.

Dave R: current algorithm in the MAC gives the same performance benefit - waiting for a NACK the same as waiting for not getting an ACK.

Rick: no, you don't have to backoff if you get a NAK.

Bob O'Hara: reduce delay for the damaged packets. What about the affect of the collisions induced by longer RTS/CTS which could increase the delay of other stations.

Rick: has no data on that, possibly.

Greg: fairness aspects of collision resolution are in doubt if not backoff when NAK.

Mike Rothenberg: perhaps a CRC protect header and use of NAK are two separate issues. Ability to more accurately assess traffic and update NAVs may be a benefit. Perhaps should be analyzed by a group.

Bob Crowder: in industrial control we have a lot of experience with NAKs. ACKs are a good thing. NAKs can lead to network saturation in high error conditions. Not sure how that would impact the current foundation. Simulation is needed. A CRC protected header of some kind may be a good independent idea.

Kerry Lynn: is it a valid assumption that all collisions will be header aligned? Couldn't a collision damage the center of a packet and then look like a CRC error.

Rick: simulation could help bring that out.

Motion #12: **That the frame format be modified to add NID, source, destination to the fixed header.**

Moved by: Rick White

Seconded by: Jon Rosdahl

Motion 12 Discussion:

Jim Schuessler: adding these fields to RTS, CTS and ACK - others have them already. Wanted to keep RTS and CTS short to minimize the collision window. Any bit added increases collision probability. The only benefit is a reduction in implementation complexity - that's what the type and control are for. State machines know how to deal with that.

Rick: simulation results last meeting showed RTS/CTS didn't have a benefit until packet lengths were large. This won't make the packets that long.

Jim S: would move further out the point at which they become beneficial.

Simon Black: benefit of RTS/CTS can be seen through ample simulations done already - they only hold if those frames are short. You are proposing adding 60 bits to CTS and 100 bits to RTS and ACK. The original Ken Biba proposal took fields out of those to make them work as well as possible. If added it might be preferable to just drop RTS/CTS entirely.

Greg Ennis: this motion might be valid as part of an entire package. But as the is motion stands by itself the addition of these fields cannot be justified.

Wim Diepstraten: functionality addressed here does not add anything new to the protocol. Function of these additions is unclear, and has large interaction with what is happening in the different PHYs. Against motion.

Dave Roberts: how many bytes added?

Rick: RTS = 6, CTS = 15, ACK = 15.

Dave R: seems like a lot of overhead, quadruples size of ACK.

Rick: doesn't like sending frames that have no addresses. Is uncomfortable with that.

Greg: this motion in isolation not a good idea. However may not be necessary for subsequent things in this proposal.

Leon Scaldeferri: without simulation couldn't decide this today. Calls the question, seconded by Greg Ennis (ayes, no nays)

Approved: 10

Opposed: 17

Abstain: 4

Motion #12 fails

Rick: straw poll who would consider protected headers (1/3, 1/3, 1/3)

unidentified: PHY and MAC need to get together here. Nothing this group seems to be discussing is independent.

Kerry: belongs in MAC/PHY interface adhoc group. correct forum. We can address it at a later time.

Rick: would it be new function to add CRC to a header?

Dave B: we talked about it before the deadline, so that's ok.

Kerry: if we beat this issue to death tonight we could be ready for the joint meeting tomorrow.

General - none of the next paper presenters are in the room, so move on.

Fragmentation Discussion

Dave B: asked for advice from PHY people last night at interface adhoc group. They immediately wanted to know under what conditions - had to say don't know. So we auctioned up from 100 bytes. Hands dropped off mostly at about 400-600. Only one was left at 1500. None above that.

Leon Scaldeferri: any PHYs looking at fragmentation themselves?

Dave B: didn't ask, but no one mentioned it.

Mark Demange: believes PHYs have discussed and passed a motion they would not do fragmentation.

Jon R: if PHY is unintelligent, knows only preamble, picks off CRC protected length, listens that duration and passes that to the MAC. MAC is responsible for what station this is destined for?

Dave B: PHY group seems to have swung direction on this a few times

Mark: Jan 93 in Agoura Hills there was a motion: PHY shall not fragment frames/packets. vote was (16,3,2)

Larry van der Jagt: PHY group was talking about this this AM. Someone in the group thinks they should, but expect a vote would come out the same again. Frame size numbers are like 400 bytes, based on 10^{-5} BER which is considered a good channel.

To stimulate conversation, Dave Bagby displays the text, which a while later becomes motion 13 - see next motion.

Dave B: better off to start simple and complicate later for time considerations.

Mark: keep in mind MAC needs to support multiple PHYs. The FH world - needs to operate efficiently. Not clear that this captures that thought process of fully supporting a FH PHY. No fragmentation loses the benefit of ensured dwell times for FH, imposes a restriction on what hop times could be used.

Pablo Brenner: network layer requirement - did we agree 600 byte was the limit?

Dave B: we had a presentation that claimed that that would satisfy a lot of upper layers. Don't know how much agreement there was.

Pablo: this motion wouldn't say what the upper number is.

Barry Dobyns: how about add 'and/or empirical data' to the end of it.

Larry vdJ: use number of 400 bytes since that's the limit came closest to being selected.

Jim Schuessler: this captures the sentiment of last night. Not sure of the value of such a motion. Isn't this going to happen anyways - if someone comes back with this later?

Dave B: if we narrow reason why something is required it will be less emotional at the time.

Greg: show of hands at the MAC/PHY interface meeting last night that reflected a split in the PHYs. If PHYs have radically different things to support the issue for us is do we have a single frame size requirement for a MAC SDU that will hold for all PHYs, or do we export the difference among PHYs up to the MAC boundary?

Dave B: so go into the middle layer?

Greg: no saying perhaps ought to be below the MAC service bound - single MSDU size for all PHYs.

Jon Rosdahl: SDU size - if small no need to fragment. But if you do, assume done in some H/W manner. Do a large SDU to accommodate that. Can be done in S/W, but there is a performance

hit. Better served to take ATM model, and come to some reason number at the SDU layer. Fragment only as need to in a manner that has no performance limitation.

Motion #13: **802.11 declines to add fragmentation function to the MAC until it is shown to be required by PHY frame length limitations or performance simulations.**

As Amended by Motion #14:

To amend motion 13 to read "802.11 agrees to add fragment function to the MAC or a convergence layer when it is shown to be required by PHY frame length limitations."

Moved by: Barry Dobyns
Seconded by: Leon Scaldeferri

Motion 13 Discussion:

Phil Belanger: 'required by frame length limitations' - how do we know if it's a limitation unless we know what is our MSDU size?

Bob Egan: this is a noop, or possibly a negative. Thought last night's discussion led to a tentative conclusion that fragmentation is necessary. This motion is not necessary, it tends to bias opinions toward fragmentation not needed. Forum for further discussion should be this evening and tomorrow in joint sessions. Strongly against. This is a motion that should be brought up as a focus of tonight's meeting. Would like to see this removed. Should be primary issue at the MAC/PHY adhoc or the joint meeting tomorrow. Request for withdrawal

Barry: declines to withdraw the motion. Thinks it performs a function in this group - it disposes of some papers while acknowledging their input and proceeds to a closure of our consideration of their input.

Bob E: dismissal of an issue that requires more discussion.

Barry: we can't decide this now, but fragmentation in a MAC shouldn't be decided in a MAC/PHY interface meeting.

Larry vdJ: there is going to be fragmentation, the physics of the real world define it. Take a packet size (400 bytes), say how many bits (3200). Probable BER of 10^{-5} , and you calculate 3.2% of all packets will be lost if no error correction done. If 3.2% lost, then retry those. This is one way to come up with the 400 byte packet size - because we are required to get 99.9% of packets through, if we say after 1 retry, if 3.2% on first try and 3.2% of the retries don't work, then you're at 99.9%. So 400 bytes seems right for trying to get success on one retry. A lot of these channels may not be that good, or sometimes better, but 10^{-5} is a pretty reasonable compromise. Another thing - radio PHYs in the ISM band have microwave oven interference bursts 4-8 msec in duration at 60 Hz (every 60 usec). Would be good to have a packet length that fit between bursts - 2, 3 or 4 msec. 400 bytes is 3.2 msec. At 10^{-3} probably a 100 byte size would be needed.

Motion #14: **To amend motion 13 to read "802.11 agrees to add fragment function to the MAC or a convergence layer when it is shown to be required by PHY frame length limitations."**

Moved by: Dave Bagby
Seconded by: Bob Egan

Motion 14 Discussion: none

Approved: ayes Opposed: no nays Abstain: -

Motion #14 passes

Motion #15: **To lay on the table motion 13**

Moved by: Bob Egan
 Seconded by: Jim Schuessler

Motion 15 Discussion: not debatable

Approved: 19 Opposed: 5 Abstain: 6 *Motion #15 passes*

Motion #13 is laid on the table

Kerry Lynn points out that technically we laid the motion on the table, that didn't lay the fragmentation discussion on the table.

Windowing discussion - agreed to postpone until after the fragmentation decision, and there wasn't one. The group generally thinks it should be left until the fragmentation decision.

Dave proposed we slide from today's agenda to next meeting - due diligence checking (no objection heard); detailed schedule review, hasn't been done, so nothing to review that is different from last meeting (no objection heard); preview of report to plenary, no time to create report, will do same as usual, is that ok (yes)

Transmit Power Control Provisions, IEEE P802.11-94/59, by Wim Diepstraten

Questions:

Kerry Lynn: dwell time on hopper, would you keep a long term average over many visits to a frequency?

Wim: keep noise level experience on operating band, not per channel.

Kerry: so it is an average of all bands a hopper visits.

Dave Bagby: one side effect is asking PHY to provide this energy level.

Wim: yes. Although not a primary need for algorithm to operate, it would be a good refinement.

Motion #16: **That the 802.11 MAC should support the necessary protocol elements to allow "tx-power control" on a per frame basis".**

Moved by: Wim Diepstraten
 Seconded by: Bob Egan

Motion 16 Discussion:

Bob E: wish simulation had been brought forward (as a general comment to everyone's presentations). May help us to resolve situations when AP has to understand who he has to repeat to - a history file could be useful for this.

Dave Bagby: there are several thing tbd in proposal. Would you bring back a paper to next meeting that filled in the unknown values?

Wim: will at least be able to put in numbers necessary to trigger discussion.

Dave B: requires AP keep some history - would like a better feel for how much information goes into this history - what it means to a H/W implementation.

Wim: algorithm that somehow averages measured attenuation, by either form of integral or maintain average over certain number of packets. Minimum would be maintained in one value per destination, maybe more.

Dave B: proposed as option? PHY group considers being told what power to use is a necessity. Don't see that you could choose not to implement it. Maybe not to use it, but we want options only when absolutely required.

Wim: by option I mean actual use of tx power control is up to the implementation. But that would mean the protocol elements that have been defined here should be mandatory to allow stations that do want to use them to be able to do so. Could be called parametric use.

Kerry Lynn: speaks against. Seems to be something that might be PHY dependent. Should be in PHY specific fields.

Wim: don't think it's PHY dependent. Tx power control benefits may differ per PHY.

Kerry: 2 of 3 may find this something they don't need.

Wim: DS and FH specifications currently say that maximum power will be one watt without power control and 100 milliwatts with power control.

Kerry: possible hazard specification to come out.

Sarosh Vesuna: likes concept. Leaves flexibility in future to do things that I may not even know yet without penalty. Thinks it is optional - if in elements field, you can build station without being able to recognize the field - if you see an element you don't recognize you just ignore it.

Wim: every station would be forced to use this element. In order to make use of this possible every station should implement a minimum - required as input in level in things sent.

Bob O'Hara: you said you want to propose a mechanism to all tx power control without corresponding algorithm to accomplish. Is not comfortable with that in a standard. Opens door for various implements to implement incompatible power control algorithms.

Wim: no compatibility issue. It is what you do in a station - you're not dependent on any other station except with respect to what they put in the frame.

Bob O: say my algorithm is tx at highest power available. Everybody else's' algorithm becomes useless.

Wim: no. A mixed environment - those that use and those don't - net result is medium reuse will increase. Beneficial for all stations.

Bob O: find that difficult to follow

Rick White: Kerry's comment about being PHY dependent - not beneficial in FH. Also variable CCA length - this is a PHY specific thing. A variable CCA, sounds along the lines of an energy detect and there has been no decision. PHYs may have different CCAs. Different transfer of information between PHY and MAC. Sounds like a PHY dependent function.

Wim: CCA mechanism can be implemented differently but the bottom line is it results in a certain sensitivity level called CCA.

Rick: don't agree, that might not be true. A lot of time spent discussing that in the PHY group.

Wim: goes back to the type of thing you have to do for CCA and the effect of the sensitivity of CCA on the total throughput of the network. Brief discussion on Monday in the tutorial that the total throughput will depend on how more sensitive your CCA threshold is compared to the data threshold (the minimum level at which data be viably received).

Rick: assume adjustable CCA, still makes sense that if PHY sets tx level and is going to modify CCA based on some vague, let phy handle it rather than passing information up to the MAC.

Wim: PHY doesn't know to what destination that frame is going. Doesn't have history needed to set tx power level appropriately per destination.

Rick: goes back to what kind of information PHY could keep.

Greg Ennis: mechanism required to support this requires knowledge of destination - this is a MAC function. Unless addressing goes in the PHY, the MAC must be involved in this. At the front of the header there is PHY adaptation - intent is to support exactly this kind of information that is appropriate for some PHYs and not for others.

Jon Rosdahl: can't use element field because it's not in the RTS, CTS or ACK. P-adaptation is a good place. To get the historical level per destination you start to place a database function in the MAC layer - how much history, for how long, and for how many destinations? Could be pretty large.

Wim: a station would probably maintain a minimum of just the AP. Could be as little as one variable. For an AP that needs similar information on a per node basis, this could be directly extended - minimum per node is one byte per address. An extension of table that already maintains power management.

Kerry: tailored to some PHYs while others don't find it useful, it is a burden some PHYs. Have you looked at using this in RTS/CTS as well - that's initial feedback between two points. Vary data level based on what you learned in the RTS/CTS exchange.

Wim: that is the idea, to use the information gained from RTS/CTS.

unidentified: in FH signal strength changes 10 to 15 dB in a matter of seconds, so you can't remember the level. It is meaningless. In FH put level in RTS, and have them all sent at high power. Then after that all frames in that transaction can come with suitable level. Then only for RTS are other networks disturbed.

Wim: this will reduce general interference level. But because of transmitting RTS at nominal levels, you have told everyone that can hear it to shut up.

unidentified: in DS, but others probably won't be on the same frequency at any given moment in FH. Best way to use power management is to allow RTS to send all the same level. You are also assuming symmetric channel.

Wim: in order to make effective you need accuracy in rx level measurement, if you send it back in the history.

Jon R: why necessary to do this power management in the data path as opposed to via the management interface. In roaming you want to increase amount of power. If managing entity is controlling, when you hit a roaming event the management could tell the radio to raise its power. People have been concerned about placing more bits in CTS and ACK. Could it go into a management path?

Wim: don't get it.

Mark Demange: power control requires up to date information. For FH the information has a very short life - might as well trash it on each hop. To Jon addressing why not in a separate path - need continuous feedback. You don't settle power level, then talk, it is dynamic.

Approved: 12

Opposed: 12

Abstain: 7

Motion #15 fails

Wim: when PHY group specifies power control requirement, how is that going to be resolved?

Dave B: don't know.

Priority in CSMA/CA to support distributed Time-Bounded Services, IEEE P802.11-94/58,
by Wim Diepstraten

Was addressed by Kerry Lynn at the last meeting.

Questions:

Kerry Lynn: if time to live parameter went below zero at source, he can discard it as well.

Wim: the time bounded LLC should do that not the MAC.

Kerry: but after it leaves the LLC, the MAC is the only one who has it.

Wim: that's why discard and time to live should be part of QOS and be carried with the frame.

Kerry: values to be able to quench traffic - could be a function of network load and the earlier you can discard the frame the better.

Motion #17: To add 'distributed time bounded' service to the found MAC.

As amended by motion 18:

To add 'distributed time bounded' service to the found MAC and to determine by the end of the May 1994 meeting whether one or both of the TBS will remain in the standard

Moved by: Wim Diepstraten

Seconded by: Kerry Lynn

Motion 17 Discussion:

Dave Bagby: might fine idea. But didn't go quite far enough in the motion. This allows a TBS that works when you go roaming. In my view that is something that is broken in the MAC because the

PAR specifies TBS and specifies roaming. So wish you has said "and replace the current tbs service with this". simpler.

Kerry Lynn: strongly in favor. Current mechanism has some difficulties, especially concerned that if users must configure for no overlapping PCFs there will be disaster. In higher rate PHYs TBS will become more important in the future.

Barry Dobyns: supports also.

Jim Schuessler: supports also as currently worded. Could not disagree more with Dave Bagby. He assume that TBS provisions only serve TBS data. This is not true in current specification - there is also asynchronous contention free traffic that can be sent in the contention free area.

Dave B: didn't say to throw away the PCF. Said don't have 2 TBS mechanisms.

Mike Fischer: supports motion as worded. Opposed to removing contention free async service - it is the most useful service available from the foundation MAC. Congratulates Wim on proposing a modification of which the costs and benefits were demonstrated in the presentation. Conceptually agrees don't want 2 TBS in every station, but it might be premature to decide there are no uses for the one that is there.

Kerry: impact of PCF TBS on distributed TBS is significant?

Wim: major.

Kerry: one of the flaws with current mechanism is - imagine a campus with labs with PCFs. One starts up, another is precluded from doing so, or second starts up and first goes to hell. Imagine a solution that is even better than PCFs. Might want to replace current TBS.

Dave B: feels stronger even than Kerry. Current TBS has a major flaw. You can't move from one BSS to another. Believes won't make it through letter ballot with current TBS service. Don't equate TBS, contention free service and PCF. Talking about TBS alone. Can't support this motion unless it also removes the other TBS. Not significant enough improvement to have the complexity of 2 services. Multiple way to do the same thing is not a standard.

Jim S: don't believe that the PAR states the need to support all modes with single and multi channel PHYs. Acceptable to have a service that works with some PHYs. Can roam using the current TBS with PHYs that have enough channels to provide continuous coverage and BSS isolation. This asks more of upper layer services than current TBS does.

Wim: we suggest here to have a non-optional distributed TBS. The current service is optional. To LLC there is a sense of connection there. Matter of translating bandwidth to maximum delay and knowing what traffic chars are.

Jim S: complex also in MAC. Have to do math to determine if frames continue.

Wim: can be left that to something above the MAC. Kerry mentioned desirable to do in MAC.

Jim S: current TBS doesn't do that.

Dave B: sees this as simpler. It is rare to run across something that is better and simpler. Can't understand why you would want both.

Motion #18: **add clause 'and to determine by the end of the May 1994 meeting whether one or both of the TBS will remain in the standard'.**

Moved by: Dave Bagby

Seconded by: Mike Fischer

Motion 18 Discussion:

Mike F: calls the question, seconded by Bob O'Hara. (ayes, no nays)

Approved: 22 Opposed: 3 Abstain: 2 **Motion #18 passes**

Motion 17 Discussion (con't):

Jim S: voted against amendment. Tempted to vote for entire motion.

Barry: existing TBS is broken in many ways, especially for single channel PHY. This offers all functional PAR requirements and would like to toss out old one in May and take this one.

Bob O: proponents of current TBS and new one should come prepared to May.

Approved: 25

Opposed: 1

Abstain: 5

Motion #17 passes

Motion #19: To add 'priority access mechanism' to the DCF of the 802.11 MAC.

Moved by: Wim Diepstraten

Seconded by: Kerry Lynn

Motion 19 Discussion:

Wim: then in my interpretation the DCF TBS is not optional.

Mike Fischer: it is not at all clear why you say they are the same because provisions of such mechanism doesn't imply aging, for instance. There are TBS services beyond just giving it a priority access mechanism to use.

Kerry Lynn: strongly in favor. Enabling a mechanism that permits DCF TBS to be constructed. At the same time there are other clients of a priority access mechanism, for example scanning for other APs. Maybe the probe response could go out with high priority.

Laura Hindy: confused - thought the PCF was optional, providing a contention free environment for async and TBS.

Dave Bagby: we invented a TBS and said it was optional even though we didn't know how to do it. Then invented contention free async which also needed PCF. Never set the balance - those two sub-things are not optional, PCF and contention free async service, only TBS is optional.

Jim Schuessler: standard says in italics (not in 93/190) "all stations must be capable of operating in an environment where a PCF is being used". They don't have to be able to convey traffic in that area, but they can't break stations that are operating in the PCF.

Laura: TBS in DCF is optional or non optional?

Dave B: DCF is non optional. TBS is optional

Barry Dobyns: calls the question, Kerry Lynn seconds. (ayes, one nay)

Approved: 17

Opposed: 2

Abstain: 9

Motion #19 passes

Motion #20: To use the mechanisms proposed in doc 58 as a basis for the "priority access method" to be included in the DCF, and to further investigate its sensitivity to relevant parameters.

Moved by: Wim Diepstraten

Seconded by: Mike Fischer

Motion 20 Discussion:

Kerry Lynn: concerned about this mechanism, would like additional information. Would like to work with you so mechanism best for the standard can be devised.

Wim: motion leaves room for further investigation. Doesn't say anything about TBS, or the QOS interface.

Point of clarification, Jeanine Valadez: adopting a priority access method - this says use this one. Does not imply anything about DCF TBS right?

Wim: yes. There is more to that than this.

Kerry: has concerns. Has done some simulation in this area, not as extensive as Wim's. Impact on default service is incurring a penalty on every single async packet. For example, low priority low load, current mechanism says to transmit when you sense the channel quiet for AIFS. This proposal adds another 16 slots to the AIFS. Another example - station about to send out a packet train of 8 responses. AIFS is currently 3 slots, so now you would wait 19 slots before first send. Then listen again for 19 slots, and so on for each packet in train. Wim assumes slot time is 23

usec. $16 \times 23 = .368$ msec. Packet itself took 4 msec, so penalty is 9.2%. Point 2 -under high load the penalties become additive. (shows simulation example) The now expanded AIFS low load penalty for every packet plus the queuing penalty. Under high load you have to pay penalty for every packet in front of you on the queue. As a percentage of payload the penalty gets greater as bit rate increases - data gets faster, slot times stay the same. Contention window is not the same for high priority traffic it's fixed, doesn't grow exponentially. Might have to limit number of stations offering high priority to avoid collisions growing. Need input from the PHY group before being able to make an intelligent decision because no idea of slot time is. If 23 usec is way to low, this could really balloon. In favor of postponing the mechanism decision until the next meeting. For those who didn't heard it last meeting I proposed an alternate mechanism: after media goes quiet, high priority stations generate a one slot time burst in the last slot of the AIFS. One slot time burden. Needs a CCA that deals with energy sensing. We again need more input from the PHY group needed and should wait until next time

Greg: separation between high priority and low ?

Kerry: high priority contention window and low priority contention window overlap, so there is a time when high and low are both contending. We don't understand the impact of this yet. My approach clearly separates them, but don't know if it's necessary yet. Wim's approach may be best, we need further study.

Chris Zegelin: trying to canvas PHY group to get slot time. Indication from FH PHY it might be as long as 80 usec. 23 sounds low for FH PHY.

Motion #21: To postpone motion to next meeting after discussion of these issues.

Moved by: Chris Zegelin

Seconded by: Jim Schuessler

Motion 21 Discussion:

Wim: specifically worded motion previous motion to allow further work to be done on Kerry's area of concern. Also showed in simulations the impact on low priority traffic. Agrees impact depends on assumptions of slot time. That's why motion says further investigate. Means that when there is a serious problem we should solve that. One way to solve it may be Kerry's method.

Kerry: like to work with Wim to reconcile simulations. Just adding queuing might change results. The impact I saw one low priority traffic was greater than he saw, would like to know why.

Wim: noticed difference - impact might be slot time assumptions.

Approved: 19

Opposed: 4

Abstain: 4

Motion #21 passes

Motion #20 postponed

MAC State machine Changes, IEEE P802.11-94/41, by Pablo Brenner

Questions:

Steve Chen: trigger time out on receive. What does rx mean - good frame with no errors?

Pablo: a complete frame. Maybe needs clarification on originals.

Chris Zegelin: since you have a slow response your station doesn't get ACK back in SIFS - you say that can be allowed for. Isn't that illegal, isn't allowing for that is unnecessary.

Pablo: have to be much more synchronized to allow for getting it a little too late or early.

Chris: no, if you don't get it in the specified time, you don't. It is an absolute.

Pablo: agrees. But you don't need a timer.

Chris: yes you do. You cannot create a machine that doesn't get a ACK out in the SIFS and be 802.11 compliant. So to allow for this in the state machine is wrong.

Pablo: philosophical disagreement, but this is just something that could be allowed because of the problem fix proposed, we don't have to do it.

Bob O'Hara: 2 issues (1) control state machine broken in 2 places: First, it goes deaf in c4 and c2 because it is waiting for a specific type of frame. If anything else rxed during that time you need to remember that and handle it. Second, if any type of frame is being rxed, the state machine needs to be notified and don't reset on the failure. (2) as a consequence of that fix it would be possible to relax timing in rxer so lower cost or S/W implementation of the timer might be possible. It could be done, it is not precluded.

Chris: FH PHY diagrams of how these machines work are very different from this. It will be interesting when you get together. When you take this diagram and couple with FH PHY diagram you will find this is broken in other ways more serious than these.

Greg Ennis: this is the rx state machine we are talking about. The critical part of timing issue is in the tx state machine. Have to be clear that ACKs need to be generated at a specific time. Must not only go out no later than a specific time but also no earlier than certain time. Not concerned about these relaxation of timer issues in this state machine.

Dave Roberts: problem was the rxer is blind to anything but what it expects to receive while it is waiting. What this change does is say if something comes to you rx it, don't just drop it because it's not what you expect. Pablo has said you might get some other benefit. Don't focus on that.

Barry: agrees its important to make this change. Extremely important. It is equally important to recognize the dependency on the size of the SIFS. We all have our own assumptions. Agrees with Pablo if we can reduce the dependency of the state machine on each timer it is a great benefit when we don't know the timers are going to be.

Motion #22: To adopt these changes and add to the draft.

As amended by motion 23

To adopt these changes and add to the draft, removing the bullet "enlarge CTS_timeout and ACK_Timeout"

Moved by: Pablo Brenner

Seconded by: Barry Dobyns

Motion 22 Discussion:

Greg: motion as stated has a bullet 'enlarge CTS_timeout and ACK_Timeout' that you said was wrong, this motion doesn't capture that.

Pablo: would remove that.

Motion #23: Amend by remove that bullet

Moved by: Bob O'Hara

Seconded by: Greg Ennis

Motion 23 Discussion:

Bob O'Hara: calls the question, seconded by Pablo Brenner (ayes, no nays)

Approved: ayes Opposed: no nays Abstain: - **Motion #23 passes**

Approved: 21 Opposed: 2 Abstain: 5 **Motion #22 passes**

General

Motion #24: To adjourn

Moved by: Jon Rosdahl

Seconded by: Jim Schuessler

March 1994

Doc: IEEE P802.11-94/82

Approved: ayes

Opposed: no nays

Abstain:

Motion #24 passes

Meeting adjourned: 5:30 PM