Tentative MAC/PHY Joint Meeting Minutes
Tuesday, January 12, 1993

The meeting was called to order by Vic Hayes at 8:35 AM. Carolyn Heide secretary.

(1) Announcements

Reception sponsored by Xircom 5-7 PM in the presidential suite.

Please give Vic computer file and printout of submissions, with correct document number, as soon as possible. If Vic gets it by next Friday it will be sent out in the first mailing.

The next interim meeting is Wilmington, May 10-13, 1993. Show of hands for those planning to attend = 27. Those that will commit to attend = 15.

September and January interim meetings are as yet un-hosted. Dates of Sept. meeting will probably be 20 to 23rd, to avoid labor day and Rosh Hashanah.

Roll call, the number of people here now means that the charge will be $78 per person.

Agenda for this meeting: submissions 93/01 and 93/07. Agenda approved by consensus.

(2) Presentation of Submissions

Mathematica Based Simulation Template for Demand Assign MAC of 92/39, P802.11-93\01,
by Larry Van der Jagt

Attempts to: (1) extend the framework presented last time to a real MAC instead of a simplistic one; and (2) explains the distribution method of the Mathematica code. The document contains encoded source, you can run it but only by paying the fee to KII do you get the source. Also distributed a price list and brochure about Mathematica. Larry also has demonstration disks off which anyone can copy files if they want them. Math source is a service at Wolfram Research, if you send electronic mail message to a server documents are sent to you. It is uuencoded and gets a bit inconvenient for large forms, but it is available.

On page 3 you see the things you get from a real MAC. The bottom of the page assigns attributes to the various A, B and C periods of this MAC protocol. Broadcast flags, message probabilities, tx and rx queues, etc. get defined and assigned. No registration process is implemented in this example. Station (STA) location is selected at random from a uniform distribution. Defines mobile STA's and fixed access points (AP's).

On page 4 there are initialization routines. The next page sets up other things. The interpreted intent of the protocol is to have one hop per frame, so parameters have to be established like hops per second, octet time, etc., which allows you to calculate the hop time. From these you can calculate the frame length, the various header times. A 5 octet frame size is chosen - short tried first.

One change has been made in the framework since last time. There were 2 or 3 states that assumed tx'ing, hard coded. Now there is a list called 'tx on states'. You supply in there the names of the states in which you are tx'ing.

On page 4 and 5 the initialization program is run on a list called AP's and a list called mobiles. On the bottom of page 5 the attenuation table is built based on distance and the statistical model. On page 6 the generate traffic model uses a Poison distribution to decide if the STA has traffic and if so puts the destination on the traffic queue.

On page 7 the actual AP state machine execution runs. Assumes no distribution system; all traffic originates from a mobile; no peer to peer C period capabilities; all traffic goes mobile-AP-mobile. The AP state machine is straight forward, moving between the defined MAC periods. The mobile state machine generates traffic, monitors SNR to determine if they can hear the AP, checks for incoming traffic during the A period which is where it constantly monitors SNR. If SNR drops to low during the A period it drops out and counts lost AP.

On page 9 the scheduler makes a reservation during the C period, randomly selecting a slot. If they make it through according to SNR they get a slot in the B period. If the AP has enough bandwidth he gives everyone what they want, if not he allocates each a portion of what they want. The algorithm for that is very simple, sharing the available bandwidth equally. This is where people will want to play around to optimize the media access control.

Larry executed it for 1000 cycles with a lot of diagnostic information. Some of the printouts are included here so you can see the output.

Wim Diepstraten: the 'down' cycle - what does that mean?
Larry: one cycle is one octet.

Execution is slow on a 16MHz 68020 MacIntosh. It needs to be optimized for performance - it is now optimized for readability. It took about 40 minutes for 1000 cycles. You pick up gross errors in the state machine quickly when you start running it.

Vic Hayes: next step?
Larry: I have already done a lot of work. If this venture meets with enough commercial success to continue I will continue, if not, I won't. Have set up a special deal for members to buy Mathematica through KII, if people buy it, it will continue.
Vic: issue list impact?
Larry: same as last time.

Federal Wireless Users Forum (Comments on NPRM 92-333), P802.11-93\07,
by Leon Scaldeferri

Mission of the forum: educate potential government wireless users about difficulties to be encountered; determine federal government needs for wireless services and record them; define issues that need to be resolved to bring wireless services into the government services; advise vendors of government application needs; facilitate government applications into standards bodies such as this; and interface to other wireless users forums.

Planning first series of workshops in May. Federal users of wireless services and people from industry to sort out issues regarding security services and other areas of interest to the federal sector. Trying to put forth requirements of federal users in an organized manner. Sees a potential of a couple of million users, and is trying to stop each group from having to develop private networks by having industry understand the unique needs of the federal government.

François Simon: do you see this to be oriented toward OSI?
Leon: not going to try to say that the ISO architecture is mandatory, but use it where it makes sense.
François: what is the involvement of NIST?
Leon: the steering committee is 8 agencies - 2 of these are NSA and NIST who have been asked to take the technical lead . Others include the Departments of Treasury, Defense, GSA, and Commerce, coming in as users.
Jim Schuessler: is your intent to keep coming back and reporting requirements to us?

Leon: yes. Also has a presentation to the MAC group about time bounded services that are special to secure voice or data over wireless, that we think are of concern to us. We hope out of the May workshop that the requirements can be put in the right language that makes sense to industry and users. Don't want federal users throwing out requirements they don't understand the impact of.
Vic Hayes: the matrix that listed 802 in some boxes - do you agree with 802.10's position there to have the services at that layer?
Leon: yes, it appears appropriate to extend some of the services that ISO stopped at layers 3 and 4. Given the nature of wireless networks some of the higher layer services need to occur lower.
Larry van der Jagt: is not too familiar with what ISO says at layer 1 for security, could you elaborate.
Leon: they talk about how you could provide confidentiality with encryption at the PHY layer. The only place traffic flow security can be done is there.
Larry: it mentions connection oriented - at the PHY layer?
Leon: PHY to PHY connections, if encrypted. The ISO document discusses how the application is made at the various layers.
Vic: is there any need for an action from 802.11 into this work?
Leon: there is one issue open on privacy. Authentication and registration were considered as part of the requirement. Some level of privacy at the 802.11 is important to us.
Larry: 802.10 isn't enough?
Leon: if 802.11 said that 802.10 sublayer above the MAC has the responsibility, to me that would be adequate. That meets the intent of the user. There are some holes for wireless, but it is well defined except for traffic flow. In a public network signalling, addressing and detection security are not a requirement. That is really the traffic flow level. The requirement for id and geo-location was requested, but that is where is the station location information must reside somewhere, and that knowledge must be protected. But the issues falling into address and signalling information and signal presence are not considered concerns.
Colin Lanzl: what does confidentiality on signalling in private network mean?
Leon: protecting numbers of the calling and called parties. May be appropriate for MAC layer.
François: the forum - has it discussed formal liaison with ISO and IEEE?
Leon: right now that is informal. There are people attending T1P1 (PCS), IEEE, and TIA groups. People attend, but no formal association providing inputs to those organizations.
Rifaat Dayem: some requirements such as location privacy and traffic flow are greater issues in WAN than LAN. In LAN the locations are limited to a local area.
Leon: you have to look at what did the user mean when he said things. Understanding the difference between WAN and LAN for instance. If your close enough to hear the emitter you probably know where he is. The ability to sit far away and access a database that says so and so is in such a hotel needs to be avoided. Understanding both what the provider means and what the user means is part of what the forum is for.

(3) Other business: none

Wednesday PM, January 13 1993

Meeting called to order at 3:40 PM, by Larry van der Jagt. Carolyn Heide secretary.

The reason for this meeting is that the PHY group has been trying to make decisions which can't be made without co-ordinating with the MAC group. Need to know what MAC should be used in thinking about what can and can't be done, and against which measurements should be made. For the frequency hopping (FH) PHY, a benchmark protocol (or protocols) is needed to determine how this PHY is going to operate.

Larry presented slides summarizing several issues which were voted on in the PHY group relative to the 2.4 G FH PHY draft specification:

1) shall require exactly one hop rate: yes-10; no-12; abstain-1.
2) shall require exactly one power level: yes-7; no-16; abstain-0.
3) shall require exactly one data rate: yes-8; no-14; abstain-1.
4) shall require tx power level control above a tbd level of transmit power: yes-15; no-6; abstain-1.
5) tbd in item (4) = (This is the level above which you shall have power control): 1 mW - 1; 10 mW - 8; 100 mW - 11; less than 1 mW - 1.
6) shall be able to transmit at least tbd level of power to conform to the standard: yes-21; no-1.
7) 1 mW is tbd in (6) 0 dbm: yes-15; no-6
8) when power control is required how many levels of power control (at least) shall we require: 4, in 2 bits, yes-9; 2, in 1 bit, yes-2; 8, in 3 bits, yes-5; 16, in 4 bits, yes-1; tbd, yes-2.
9) shall hop rate be configurable in the MAC but fixed for a given BSA if it does not have to adapt: yes-20; no-1; abstain-1.
10) MAC will tell the PHY when to hop: yes-10; no-12. voted again and got yes-13; no-5
11) the PHY shall not fragment packets supplied by the MAC: yes-16; no-3; abstain-2.
12) shall require the MAC to maximize the use of each hop interval: yes-15; no-4; abstain-3.

Then a slide showing where things get MAC/PHY-ish. "In hop" indication needs to go from PHY to MAC; "Time to hop end" needs to go from PHY to MAC; and "hop request" needs to go from MAC to PHY. The error rate has an impact on the hop rate, and the channel performance. PHY group decided there will be more than one hop rate, and it should be something that is set by the MAC, and there is no intention of that to be changeable according to conditions.

Discussion:
Dave Bagby: static durations? We won't change the length of each hop? What's the dwell time after each hop?
Larry: hop rate is fixed. Some mechanism in the MAC causes things to recognize the hop rate is such and such and they are fixed sized hops.
Wim Diepstraten: the PHY doesn't change this itself - it is under control of the MAC?
Larry: the next question is - yes it is under control of the MAC but we would like to participate, share in the process of determining this. Let's just talk about the hop problems. Fragmenting - a thing that is under some error control mechanism like a CRC you want to make sure that that thing that goes out doesn't get fragmented over two hops. When you hop that's when you get into differing error conditions. The PHY won't fragment packets from the MAC. How does the MAC know when the PHY hops? We thought the MAC would tell the PHY when to hop - we can't come up with any other option if the packets are not going to be fragmented between hops.
Dave B: no buffer in between - it would be nice if you could hand a packet to the PHY and have it take care of it.
Larry: then we could wind up sending it when the media conditions are no longer valid, if we couldn't send it when MAC told us.
Colin Lanzl: we discussed this in the PHY group - if the MAC is in control and telling the PHY when to hop, then we don't know how to do that. Some control function is required, but we don't know what.
Larry: there must be something providing a global BSA synchronization clock.
Colin: think about a portable going from BSA to BSA, in the case of non-centralized protocols it is not clear how this synchronization is maintained.
Tom Baumgartner: if you have fixed hop periods and you don't want to break data across hop periods, then either the MAC knows how many bits you can get out and divides data properly, or time will be wasted waiting for time to next hop.
Larry: in terms of a real MAC, if there is a superframe allocated into various uses - if that frame is a hop, the global MAC thing decides how to use the time.

Chandos Rypinski: it is difficult to develop a MAC which is medium independent. There is a lot for a MAC to do besides managing channel or hops. In the alternate, how about a PHY service layer which has the functions which are PHY dependent and the MAC can stay standardized. This may be better partitioning of functions.
Larry: if the MAC decides now is the time to use the media, and there is something in between that decides if you can use the media now, then that invalidates the MAC's decision to use the media.
Steve Chen: PHY can say that this is the time period available to MAC and if MAC decides that that time is no good, then it waits.
Dave B: the upper end of MAC gets data asynchronously from upper layers and at the bottom is a PHY with very periodic behavior. At some point in that flow "now is the time it can be sent" has to be determined. Whether MAC or a middle layer does it may not matter. What you want is a continuous stable/not stable/stable/... information, then can you get at the information quickly enough to be useful. I might tend to call that a middle layer because it is for a periodic FH PHY only.
Larry: you are looking for CTS?
Dave B: yes, but can we respond fast enough to use it, and will it be ready long enough. Will it still be ready at the end of the data.
Steve: PHY can tell MAC "this is the remaining time of the period available", then MAC can decide to use it or not according to the access protocol.
Dave B: thinks of a chunk of buffer in between loaded up and ready to go, and when the hardware was ready it got dumped. That seems simple.
Wim: thinks MAC needs control of the PHY because MAC should have a timely exchange of information to perform that access function which may be time critical. The amount of time left in the hop is very important. If the MAC is changing the hop, then it has the knowledge to calculate that. There are issues of how, but the control belongs in the MAC. Don't add a lot of functionality just to support a FH PHY.
Wayne Moyers: there is every reason to expect some asynchronization between the MAC having something to pass and the PHY being quite ready. You have to have a buffer, bounded by packet dimensions. Then the PHY sends it when ready.
Dave B: MAC has to be in control of PHY but doesn't want to be constantly looking at the PHY to determine conditions. Load it with parameters such as hop frequency and length and say go. Some sort of and gate so that you can tell when there's data in the buffer and maybe some other conditions. Then the same sort of thing other way on the receive side.
Larry: think about the timing of this stuff - the fastest we are talking about hopping is 25 ms, so you are inside of a hop for a long time.
Dave B: you said that PHY will not fragment a packet - that is very nice to hear because it helps encryption and security stuff. What is it that you think is a reasonable expectation for channel reliability?
Larry: personally, a bit error rate of 10^{-3}, others feel 10^{-5}. Uniformly we agree on the idea that a channel is good or bad. If it's bad it is probably bad for the whole hop. The most likely way to get better performance is to go to the next hop. There is little sense in re-trying in the same hop.
Mike Bergman: what is the dead time between the end of a decent packet when you hop and the beginning of the next packet. Probably for isochronous data we should limit delay to 4 ms. Also, why not fragment packets in the PHY?
Larry: Let's say - 300 µs to hop. That's 30 octets. 2.5 hops/sec = 400 ms, which is 40,000 octets. At 40 hops/sec = 25 ms/hop = 2500 octets. Plus some maybe 100 bit preamble. That's the kind of scenario you can expect. If you transmit a packet and it gets bombed, the best chance to get it through is to wait another 40,000 octets and try again on the next hop. If we fragment then you don't know what hop the problem occurred on. We want to transmit what you give us all on one channel, as characteristics are different on other.
Tom Tsoulogiannis: you are assuming packet loss due to error, not collision. The retry mechanism is at MAC level - how does MAC control it if retry is in next hop? Assuming we

were doing only FH PHY, would this function arise - it is only because the MAC is PHY independent that we don't just say have the MAC do it. Don't know if a MAC totally independent of PHY is possible given this FH stuff.

Bob Buaas: take the example case of MAC doing its protocol decides to transmit. The PHY layer needs to change frequency and so it doesn't get the MAC's first bit out for another 300 µs. The MAC must have this in mind - assuming a particular PHY implementation. These are the kinds of questions we must consider.

Nathan Silberman: we are complicating things unnecessary they have been resolved in other places. They have not been explored well enough yet. Things presented here are a result of a binary decision tree followed in the PHY group. Also 300 µs is unrealistic, 100 µs is more probable.

Dave B: 300, 100 doesn't matter call it x and y if you want.

Tom T: when I decide to transmit I have x bytes to go. the PHY cannot tell unless it knows the length if it can get the transmit done in the hop window. The MAC has to know where the hop interval finishes. Again the MAC becomes PHY dependent.

Dave B: you can isolate these functions to some thin layer.

Colin: a station should transmit to a station and expect an ack back to verify the channel. If you do that then and it works then you want to stay with that channel. If you have more to go than can fit in that hop it starts to get hairy.

Bob B: when you see the PHY group report you will see that the voting on the questions will show the threads we followed and what the feeling was. Some questions were balanced and we tried to look at both (or all) paths. The vote numbers you may find useful when balanced against the questions asked.

Dave B: in a FH you have to synchronize the devices. How does that cross the boundary?

Larry: we haven't had enough time to address it. That's almost why I say we have to assume a MAC before we can do much. If there is no co-ordination function (CF) this PHY won't be able to sync (some PHY group member maintained). Some kind of MAC assumption would help to see under x set of assumptions how will it behave. The sync problem for instance could be totally a MAC problem.

Dave B: once the devices are sync'ed it might as well be a non-hopper except that my ability to use the channel is not continuous.

Greg: another MAC mechanism may be impacted. If some request/response mechanism is required it may have a timing constraint, and we are looking at a PHY level where the time between when the MAC sends the last bit to the PHY and when that gets to the other station - that may vary by at least 300 µs. A buffer between the 2 for instance, could cause problems if the MAC doesn't get to control exactly when data gets sent.

Dave B: when a certain amount of time elapses I have to hop, ready or not.

Bob B: what about where some portable may be sleeping, and has to try to estimate when to come back to life - if its variable you don't know when to wake up.

Colin: if we assume peer to peer and the hop control is resident on one MAC - that may be too much for that one MAC to do.

Dave B: we've admitted that possibility and don't know how to do it either.

Wim: it is clear that the MAC should have control over all. The next challenge is how to partition into PHY dependent and PHY independent and whether we can simply parameters that.

Larry: if the MAC tells the PHY when to hop, its telling PHY change channels. PHY doesn't know it's FH, it just obeys a change channel command.

Wim: we have a different understanding of what CF means. CF to MAC means rules for how you get access to the network, while you mean some co-ordination to synchronize the hop sequence.

Larry: how you access the network and when, and on what channel is in essence your cf.

Leon Scaldeferri: ad-hoc network is impacted by this. How do you maintain sync and allow ad-hoc and non-ad-hoc to coexist? Someone has to be in control of that hopping, and who takes that role is a serious problem.

Larry: if there is a possibility of a MAC without a CF (i.e. ad-hoc) we are worried.
Dave B: basically, you want your timing queues to come from something external to the PHY.
Chan: it is possible to sync all these guys without too much power. Everyone has a clock - how long is it before the difference between those clocks exceeds a tolerance. One needs to beat the drum every tolerance period. You may be able to go 10 minutes without hitting the drum.
Steve: don't think telling when to hop is a MAC function. If we are required to have one MAC we can't do this there. These functions for direct sequence or infrared are not necessary in the MAC.
Larry: the real world dictates where these things will be - if we can't solve the problem we have no MAC instead of one. We may have to change our requirements.
Don: if every station is sleeping in a peer to peer situation - if some one wants to come on before the end of the hop occurs, you don't want to wait for that 10 minutes.
Larry: someone has to tell everybody in a BSA what channel to be on at what time.
Dave B: but I could load you up with a sequence and you listen until you hear someone else in the sequence and join him when you know.
Colin: you may have a similar function in a direct sequence - if you decide to change spreading sequence you need to tell everyone. Parts of this may not be unique to FH.
Bob B: agrees with Colin. Some one has to be in charge of time. It is important to several MAC functions, so it makes sense to centralize co-ordination of time in one place.
Chan: stepping pattern is algorithmic and only an initialization of clocks is it necessary, and it depends on the drift of clocks in stations relative to each other.
Steve: if we have one single MAC - maybe w should consider architecture and reconsider a station management block that goes across the MAC and PHY parallel to both and handles PHY unique functions, but is not part of MAC. The two groups want each other to handle it.
Nathan S: Any radio PHY requires some sort of sync or ticking clock somewhere. In FH it is hop, in direct sequence it is called something else. If extended TDMA used, for instance, it needs clock too. Distribution of timing is essential for all.
Dave B: there is a MAC and many PHYs and some udder shaped thing between them in the model we have drawn before. I don't care if you call this middle thing MAC or PHY. Does PHY still feel that previous ideas about relative intelligence level still hold? There was a desire to make PHY unintelligent, but maybe the middle layer is PHY and intelligence is there.
Larry: we thought we could upload some thing to the MAC - any information protected by CRC has to come from MAC and all functions related to that.
Wayne: the MAC could tell the PHY - end of hop.
Larry: if there is a superframe the hop can be exactly the size of the hop. Then the MAC knows exactly how long the hop is.
Dave B: the MAC proposal by KS Natarajan has that assumption, but Ken Biba's proposal is much more contention based and doesn't allow that. Maybe these are the two you should chose because they are so extreme in this area. When you talk about when to listen and when to talk that is what MAC thinks of as cf. You would like dwell time information from us, start and end, these are different from when to talk and when to listen.
Larry: you could tell us start a sequence and we could go run with it. But in light of sleep time, retry time, ... these issues come into play in that decision.
Dave B: if we just started PHY off, some information would have to come back up, like are you dwelling or hopping now, etc.
Larry: don't think you should send data if the PHY is not ready for it. You could just wait until you get the right indication from PHY before you give it data. PHY would have to stay awake tracking the hopping sequence.
Wim: we may find out when we are looking into power management requirement that we need similar timing synchronization and similar awareness of time in the MAC, and building these extra hooks into the MAC to allow these kind of FH support functions is then only a small addition.

Meeting adjourned: 5 PM.