MAC Minutes
Wednesday July 8, 1992

In the absence of chairman Dave Bagby, the meeting was called to order by Simon Black at 8:45 AM. Carolyn Heide secretary. Simon would like to continue with MAC/PHY Interface because we were making such good progress yesterday in the joint MAC/PHY meeting.

Simon puts up a diagram as a discussion start, not gosple. [sec note: the following diagram is what we had on display by the *end* of the dicussion.]

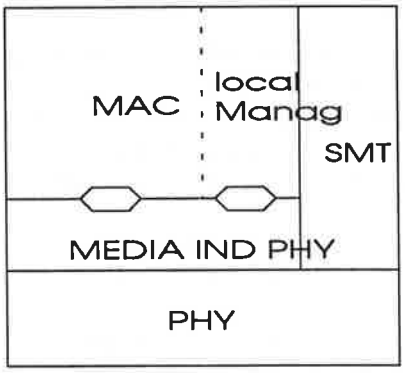

Discussion:

Paul Eastman: this is not a matter of whether there is a management function or not - one has to exist, but do we seperate the two - integrate into the MAC or not.

Simon Black: this has much to do with what conceptual model you use. We need to try and clarify that.

Tom Karota: MAC vs mangagement - if you define MAC as anything that's not PHY then it's part of this group's job to define it. Whether on the side of the MAC or not is just tryng to define it too well - the management agent does not need to be seperate. MAC is just not anything PHY, so it's MAC.

Andy Luque: reasons for putting it into the MAC - seperate the 2 things. Logical definition only, doesn't guide implementation. Identify management logically, this opens up remote management. A central agent rather than considering knobs to twiddle is clearer, more concise. All in one group - other 802 groups have management done in the MAC group, but logical seperation is good at a logical level.

Francois Simon: I support that view also.

Paul Congdon: traditional management involves remote things too. Packet by packet basis things may not be considered management.

Simon: Larry said some things may have to be adjusted on a packet basis, some may not.

Jonathon Cheah: the PHY I know requires management on a bit by bit basis. Regular PHY requires control bits to tell you how good an octet is. In a serious digital radio situation each bit is managed. That is the basis that causes a lot of PHY guys to get hot under the collar. I have a problem with the management structure as an independent layer. Now the PHY has a MAC and a management interface - that doesn't sit well with the PHY people. PHY gives a primitive with a control word, MAC passes back a control word which encompasses all the things you want to know for these n-bits.

Andy: you want to distinguish between the bit or fragment and the long term management. Remember that some of the information stored - quality of service sort of thing - turns short term information into long term information.

Jonathon: immediate management and long term management are seperate.

Jim Schuessler: and on different interface's.
Bob Crowder: idea of control words is what I was describing - bit by bit is a little frequent. In the OSI model for management there are 5 categories of management. Redundancy or media management can easily act on a frame by frame basis in parallel with the MAC, but still seperate. So you have a very specialized MAC management that has a PHY concept built into it - in 5 years that algorithm will be obsolete. If we could define something very quickly that all PHYs could use. How to handle media that requires special functions could not be left to the side, but it could be left out of the main line of the MAC. It is easier to deal with new primitives than new states in your state machine.
Jonathon: I agree with Bob. Controls should be on primitives. Back to bit by bit - the MAC doesn't need to manage bit by bit, but if we need FEC (forward error correction) does it go in the MAC or PHY? If MAC refuses to manage bit by bit - send a message to PHY that says if you do FEC do it yourself. In the PHY group people want to make the PHY as simple as possible - if we have plug replaceable PHY then it makes more sense to have FEC in the MAC since we haven't decided if we want it or how to do it. We can say bit by bit is no good - we can make that decision and tell the PHY group.
Bob C.: FEC is not what I meant by management. Media management to me means antenna selection, frequency hop pattern selection, etc. It doesn't make sense to change these thing mid-packet.
Simon: in field bus, does the media management information - control words frame by frame management ... do those words flow across the MAC to media independent layer interface? Do they actually flow to the side box?
Bob C.: yes.
Simon: two types of management envisioned - typical LAN management - status, observation and control information typical of Ethernet and Token Ring; and the frame by frame basis management information via primitives that flow between the MAC and PHY interface - data and control primitives going across MAC/PHY interface. Then how far does the management box go down the side?
Jonathon: PHY wants to see a contiguous interface. PHY doesn't want 2 interfaces - one to the MAC and one to management.
Wim Diepstraten: what about the traditional management interface?
Jonathon: what is traditional management - phase lock loop? Control primitives all go into MAC. MAC controls frequency - everything that the PHY does. There is no management in the PHY.
Jim: there is a path between the SMT and the PHY always traditionally.
Jonathon: as far as PHY is concerned the line between PHY and MAC is artibtrary.
Simon: logical interface is a service access point on the interface.
Paul Eastman: In 802.3, 802.4 and 802.5 all the management takes place at MAC and above. PHY does not.
Andy: you cannot address the PHY, independent of the MAC, remotely.
Simon: in OSI management covers all 7 layers. We can't assume that there is information in the PHY as managed objects that we have no knowledge of. The object has a history of the media access - MAC should be able to get any information, so the layer must go all the way down.
Jim: agrees with Simon, it is mandatory. In FDDI the management layer goes all the way down and PHY has management objects. There are things that the PHY knows about that are unique to the PHY - link error montioring is done based on those. The PHY should take care of those and the SMT should get that information.
Bob C.: agrees with Jim. Main reason for this solution is to provide an explicit management interface to PHY with MAC. It works like Jonathon said, except that the independence is achieved in the PHY independent layer - although part of the

PHY, it is implemented by MAC guys. Through that interface normal get/set primitives can be passed. How did this multi-level scheme come up - about 7 pins can get all the traditional management and select antenna and tx data. The real PHY - below the DTE/DCE - only needs those 7 pins. It allows media PHY and MAC management very simply.

Jonathon: I think I agree with Bob. The biggest problem we have is not identifying which part of the PHY independent layer belongs to MAC. DTE/DCE interface becomes a logical break. This independent thin layer is neither PHY nor management. The management comes down to the DTE/DCE seperation point - then I have no problem. Then the bit by bit management resides in the independent layer - those statistics will come out of the error counts that reside in that layer. Since the MAC group is implementing this layer then the management can just come in - they all sit in the same silicon with the PHY.

Mike Bergman: sees a need for 2 pieces of management. Local for things specific to the unit, and a traditional management function. The latter should be a block along the side of the MAC and PHY. 2 functions, one for local and one for statistics. Which pieice resides in which box becomes the issue.

Jonathon: redundancy is a multi-PHY situation, where MAC has choice of antenna. This is a bad name - redundancy means coding bits, tx media options perhaps.

Paul C.: PHY layer loopback, PHY reset are also types of things that come down from MAC. PHY shut off for jabber control.

Wim: 2 types of management functions - in MAC and traditional SMT only down to the PHY independent layer. In the traditional 7 layer management there is a management block and then inside the MAC there is another block. Rather than going with an extra side block - call the block on the original drawing the traditional SMT. Also have a seperate block inside the MAC block (next to the SMT).

Simon: we need a service access point for OSI reporting and control primitives on the side of the PHY independent layer, and SAPs at the interface between MAC and the local management and the PHY independent layer.

Andy: maybe a dashed line between SMT and MAC local management - a soft interface.

Francois: my original foil looks just like this drawing.

Andy: problem with this is the local management and OSI management need to be tight - you don't want a formal interface.

Simon: it is only logical.

Andy: once you are doing request/confirm it is formal. This interface needs to very tight.

Jonathon: it could be implmented by DMA ...

Andy: request/indication puts constraints on the way the information goes across, you can't just put your hand in and grab something. There is a dotted relationship - not just from the point of view of implementation.

Wim: the proposal is that the traditional management could go all the way down to PHY.

Andy: the issue of whether you go all the way to PHY is different. The formal interface problem is that if OSI wants to increment a counter a event must be generated, etc ... why burden yourself. If they both sleep in the same bed then they won't have to pass information across the night stand.

Wim: it is at implementation that these interfaces get merged.

Andy: even at the logical level it imposes a significant additional burden at that level of formally.

Francois: on the text of the standard.

Nathan Tobol: the way something is described has nothing to do with the way its implemented. Describing things in very formal terms - clean and unambiguous with

request and response and primitivies - provides a good clean interface. Implemation is outside the scope of this group and can be done by just getting in there and mucking with the counters. If the interface is described cleanly venders can implement things proprietarily. Using the primitive describes exactly what happens - then each implemention will function the same.
Andy: we may be violently agreeing. Key distinctions: local management thing is related to doing frame by frame management. Other question is whether to count frames there too - so there is a gray area as to what fits into which box. If it is doe in the local management box then the formal interface is ok because all interface is through true OSI components. That goes out to other entities - I have no problem with that - it is a question of what goes on which side of that interface.
Simon: we don't know that yet.
Andy: the SMT needs to be truely OSI management and the local one truely 802.11 things.
Jonathon: Francois's and Simon's diagrams are now the same. Andy's analysis is good. It takes issue on Francois's PHY box - the PHY on Francois's diagram is the media independent PHY and below it is the DTE/DCE PHY layer. The SMT only goes down to the media independent PHY.
Wim: buts that's all PHY.
Francois: the SMT box goes down the whole PHY.
Jonathon: but most of the PHY the SMT does not see.
Francois: all counters maintatined by PHY will be above the DTE/DCE layer?
Jonathon: yes. There is nothing down there that you need to know.
Simon: all managed things in the PHY are media indep?
Andy: good idea - solves the problem of PHY not knowing addresses.
Simon: we have agreement on at least a model. Specifying objects - maybe we need PHY group feedback.
Mike: bits, frames or octets across the interface - can we discuss that too. Agrees with Bob Crowder that frames is less burdensome - bits or octets is over optimizing.
Simon: we are suggesting control on frame by frame basis at the service interface at the MAC to PHY independent layer.
Jonathon: too slow.
Bob C.: how could you change antenna in mid-frame?
Jonathon: don't want to store data frames in the PHY.
Bob C.: issue is when can the MAC changes antenna or frequency hopping pattern. Between frames, not in the middle of them.
unidentified: there may be fast frequency hop that hops multiple times per frame.
Simon: those will be controlled by PHY on MAC instruction. You would set the pattern for the PHY on initialization. A list of frequency hops given to PHY on initialization.
Wim: perhaps a frequency hopper could be a test for how this works. There are issues where there could be problems, like who is switching frequencies.
Jonathon: who does what is the question here. There is a a gray area where you push the power curve up or down. Could buy Bob's arguement that if you have the set of frequencies at the start of the frame, then at the end of the frame a result - there would have to be an agreement with the other end first too. The MAC would tell the PHY - I'm no good, you optimize it. That control gets done in the PHY independent layer (PIL). The PIL then tries to do something for you. I don't like that - a frame is just too long, now the PIL has to make very complicated decisions which require interface with outside - would like to push this up to MAC.
Simon: but if you do that you've made the MAC dependent on the PHY. The MAC doesn't know about frequency hopping.

Jonathon: PHY layer options - frequency hop lists is no different from a spreading code.
Wim: there is a synchronization issue within a hopping sequence. Who decides when the switch is done?
Jonathon: has to be MAC.
Chandos Rypinski: has been suffering through this intense dicussion of frequency hopping (FH). In my view FH won't be what we settle on - occupied spectrum as a ratio to tx rate is going to be too large. 75 or 79 channels have 1/79 of the occupied channel. Not many will go that way. So the administration of FH tables should not be considered ad nosium. As far as channelization and the difference between spreading and FH there is an entirely different relationship of time constants. Headache - channelization be it frequency, time or code it is still channelization and we need to consider channel selection as a MAC function and those results need to transmitted to the PHY. Can concieve of selecting channel more than once per frame - everything is possible but not necessarily preferable. Multiple antenna - take the case of a portable computer with 2 antenna since there may be needed to get 360 degree coverage. So there is a channel selection - this is a PHY selection not a MAC problem. There are MAC selection cases which are AP problems not station problems.
Jonathon: what PHY guys are going to do - we have agreed that they would give us a set of primitives that transcends - a set of universal information that allows MAC to make the right decisions. To determine whether MAC makes decisions in a frame manner we have to wait and see what the set of primitves are. We have to work with what they give us. It is premature to decide frame by frame or subframe now.
Bob C.: what can you do in the middle of a frame the doesn't destroy that frame?
Jonathon: optimization is a continuous process. If you are getting 2 signals from 2 antennae, and selecting the best signal between the two - before the packet is ending I would have done some optimization. There are functions like that.
Bob C.: how was the MAC involved?
Jonathon: can say you are allowed to optimize.
Chan: diversity combining goes with the possiblity of multiple antenna. A diversity combiner can make a decision in the middle of the bit - it is a physical media problem. The MAC should not have anything to do with this decision. The simplest form at the media level is seeing one signal is good and ignoring the other. Hooks can be put in for this interface as needed. This is a PHY function.
Wim: agrees with Chan. There are PHY decisons to make without the MAC and this is one. A receiver selecting the best antenna is a PHY function. In answer to Bob Crowder's question - is there an example for management within a frame? In a FH environment a large frame is going out - the switching to another frequency may be scheduled within that frame. When the MAC doesn't know the timing of what is going on what happens to that frame? Does the PHY say wait for a while?
Jim: flow control between the PHY and the MAC will take care of that.
Jonathon: Chan argued for me - how would you be able to tell diversity? PHY does not read the data. That is why anything requiring reading the data has be MAC and MAC tells PHY what to do. Another problem - MAC frame might not be the same as PHY frame. Burst or continuous flow - PHY might work on a different frame struture than MAC. Don't close the door at frame level - open it at subframe.
Chan: flow control - the problem with FH is there has to be a re-acquisition. If there is consecutive transmission on 2 different frequencies then there is a frame for each frequency so that initiailization and acquisition is probived for. Selecting between 2 antennas - there are multiple ways - we don't know that a packet is good until the MAC has read the CRC. But you can tell a bad packet in the PHY. FEC should and must be in the PHY because the kind of FEC is media dependent. All FECs have

consideration about current drain and delay time trade off - Solomon decding is going to take time and battery drain. The MAC should not be involved in this - each medium has its own optimizations. Selecting diversity - you have only one source for the signal but you need to select. Simplest is once per frame. Reluctant to change sources in mid-bitstream particularly at higher rates. Power and analog combination are no way a MAC concern. All we have to decide now is whether antenna selection can be done at the PHY without MAC.

Simon: summary: MAC can request some status or control information on a frame by frame basis because MAC deals in frames - on receive doesn't know if a frame is ok until the CRC has been checked. But PHY - diversity frequency, space, or antenna - might be done by the PHY on a bit by bit basis independent of MAC. Information across the interface is only frame information implemented in the PHY not mid-frame.

Jonathon: no. If one MAC and multiple PHYs, the frame structure of the PHY may not be the same as that of MAC. So PHY might have to make changes mid-frame.

Chan: only FH.

Jonathon: IR - you could have a damn long frame.

Jim: how does that affect the MAC - say you have a long MAC and a short PHY frame. How does that affect the MAC?

Jonathon: PHY can't read data. MAC must sit waiting until PHY has finsihed transmitting.

Simon: MAC receives frames from PHY - MAC can't know what was no good until completely received. Can't know until end of frame.

Jonathon: MAC decided at beginning of control word of preamble to make sub-frame decision. In local management reading data needs to be able to react to different frame stuctures in PHY.

Bob C.: hopes that PHY can be configured to not hop in the middle of a frame. MAC configures PHY to "no hop during frames". This must be done to have a generic MAC. Can't have fourteen different states because FH in progress. You have pointed out information that comes up during frames - that is bit by bit information to which "I think this is an error" is a valid response and the normal reaction is to trash the frame. That information goes up the left side of the diagram.

Jonathon: MAC frame is the minimum of all possibe PHYs.

Bob C.: you could configure MAC for different frame lengths for each PHY, but minimum MAC = minimum PHY. But don't penalize none - FH with small frames.

Chan: another issue - distributed transaction and one at time protocol and MAC. One at a time is where all transmissions between 2 stations have no interuptions. If time slotted protocol and some slot is control what you have is plural transactions simultaneously within different stations. The stations have to react accordingly. This is a disadvantage - any time synchronization between request/grant and use required is a bad thing because of the complex states that exist in a pending condition, control requests come from concurrent parallal tranactions. We should avoid this.

Ian Crayford: we are hung up on optimizations in MAC - MAC is a frame based device. PHY is a bit based device. Optimizations on receive, the MAC doesn't care what part of the frame is bad - the frame is bad.

Simon: but you could opt to avoid things getting bad in the frame.

Ian: but PHY knows that, not MAC. MAC/PHY frame size need not be the same. There could be primitives at reset that the PHY says to MAC "my frame size will be xx" and the MAC adapts to that. Not a frame by frame thing - an operational parameter.

Mike: the MAC can get and set parameters within the PHY and it reveives indication that the PHY has made the change independent of the MAC. Also the

MAC could lock and unlock an option within the PHY - so that the MAC may only lock and unlock some things that the PHY can change independently.

Break from 10:10 to 10:45

Simon summarizes:

MAC receives information about the PHY and can program PHY on reset (for frame size, available 'controls', etc.) or on request, not reset.

MAC can request to change PHY parameters on a frame by frame basis. What can be changed is to be suplied by PHY group - probably diversity options, power control.

PHY can internally act to optimize on a bitwise basis. MAC may be informed of what changed, but has no control.

Now let's try to define service and service parameters

Discussion:

unidentified: in terms of standards documents, what is the next level of document after the functional requirements.

Simon: I think an architecture document is next according to document plan for the 802.11 group.

Francois Simon: are we going to attempt to close some issues?

Simon: attempt to close issues when the groups are back together again. We could open a lot of new issues now. Does anyone have any now?

John Deane: yes -

How does PHY power-down affect MAC?

Alternatives: 1. re-register at turn-on; 2. coordinate turn on with AP.

Pro: 1. registration function has to exist anyhow. 2. stations behave predictably - AP can hold store and forward MSDUs.

Con: AP doesn't know if the station has gone - ie when.

Jim Schuessler: PHY can power down independent of MAC control (or under MAC control) so the question is how to coordinate.

John D.: ad-hoc networks are a particular problem.

Chandos Rypinski: I have considered this in my proposal - polling is essential for a lot of reasons, one specific example is for sleep or not sleep. If polling shall not take more than 1/2 the time (one poll each second or 1/2 second) it is possible to specify out of MAC how long PHY shall sleep until wake again. This is an initial proposal and don't know if it's right yet. But the control must come out of the MAC and will eventually get down to specified parameters. Sleep mode should be a command set. Can sleep mode have something to do with the reciever sleeping until a signal comes in - receiver listening 100% of the time is a significant source of battery drain. That and the possibilty of channel busy continuously - sleep mode from listen as a PHY autonomous procedure won't work.

Jim: are alternatives 1 and 2 mutually exclusive?

John D.: station tells AP when it has come back to life - not necessarily exlusive, but that might be simpler. Pretend a station is in communication with an AP. It is given a poll and says I'll sleep for x time. The AP buffers its messages for that length of time and when it's told the station has woken up it sends its stuff down.

Jonathon Cheah: simpler approach - turn on/off willing or forced. Think about a very secure environment where a message is sent but intended receiver hasn't received it

- how important is it to keep this message? If the 2 parties are not right there, the message has no value. If assume that, and then if a station is off, the AP sends 3 times then discards the message for him and replies to all subsequent messages that the station is not there. OR assume station is always listening and while the network is going you can hear it. You don't have to register, you just say you want to send to a station and AP says he's not there.
John D.: authentication?
Jonathon: if station requests to send a packet to station', AP does not give ccess to station' if station does not give the right password. All this is possible if you assume that if the two willing parties are not there you can't consumate.
John D.: you could hold messages because the AP knows that the station will be back because it is just sleeping.
Jonathon: you are proposing something slightly higher level. E-mail sort of holding messages until the party is there. You relay information upward that the station is not there, or just return up that the station is not there. The onus of sending should be with the higher level sender.
John D: disagrees, there can be robustness improved.
Ian Crayford: registration and authentication - aren't they way outside 802.11? The network looks after this for you. If you do something else you will have problem with exitsing networks that do this already. The hub is a PHY layer bit repeating device. Security, management these features are outside the hub functions. We shouldn't be addressing these issues - infrastucture issues are not our responsibilty.
Jim: agrees, this is above the MAC layer, but these are unique wireless issues. We have to handle them and that is why the ESS/BSS concept exists. Because existing operating systems don't know how handle mobility we must do it.
Ian: they are doing this rght now (Novell, LAN Manager, etc). We can do link layer management, we shouldn't do session layer. Nodes going away, intermitent nodes, are handled by existing networks.
Jim: the pro for alternative 1 is that its simpler. The con - a simple timeout would suffice.
Ian: there are a number of proposals and not all have addressed this.
Jim: this is a very valid issue.
Chan: registration and sign-on are 2 different things to coordinate security. When you say registration you are talking about sign on.
Ian: yes.
Chan: it is intolerable for an AP or infrastructure not to know who is there and where they are before there is a message for them. You cannot find out at the instance the message arrives - use the idle time to get that information. All you need to know about a station you have to accumulate in off time (power, antenna, etc).
Ian: the power control issue is very important to network technology and needs to be handled as a managemented object and administered. Have to be able to write to the device at all times - administered from MAC to PHY or vice versa, or an external device administers to both.
Simon: power control can go down the SMT stack. Sign-on, registration, is an above the MAC issue. Power control is a relevent MAC issue.
Jonathon: the rx part and the tx part are different as far as registration goes. Authenication is not the problem (it is someone else's responsibilty), how to do sign-on is our responsibility. The functional issue is do we want a network with no memory, some memory, or humungous memory.
Ian: provide hooks but allow higher levels to provide that stuff. Ethernet networks don't register every id - they come up and some else handles that. We have special considerations, but looking inside packets is not our business.

Jonathon: no memory system - much simpler and allows a lot of options. Frees how the wireless network can be used with multiple PHYs. The essence of a memory-less system is the sender takes the onous of everything. I assume that the world does not remember what I sent. I send a packet, if it is not acked within a time, I assume the packet got lost and is forgotten, and I must take action if I want to. The sender has all the information to make that decison memory-less. Infinite memory is required for guarenteed delivery service - we don't want this. Our information is timely, it expires.
Wim: is this subject relevent to what we are or were doing right now? These are important issues - registration, sign-on done in MAC or higher layer, for instance - but maybe there should be a seperate time for that. Does this issue have an implication on what we were talking about?
Simon: what do you think?
Wim: the MAC controls PHY power down. It is not a PHY function.
Mike Bergman: it is a PHY function if a re-sign-on signal is used. When no traffic, continuously re-sign-on to keep track of the situation.
Simon: you're assuming infrastruture, this is no good in ad-hoc networks.
unidentified: power down affects MAC or not? In 802.4 they still had to deal with this in the state machine - offline and online is similar. If we rely on higher levels it may be distruptive. Turning off a PC should not affect the whole network. MAC should have its own management function aside from uper level functions - they decided this in 802.4, not to burden the upper layers.
Simon: to turn of the receiver for power saving is one thing. The other is what registration/authentication function knows when you turn your machine off and take it away. What we are dealing with is affected by this when the receiver turns off to save power - MAC puts PHY in sleep mode or at least has to know when PHY went into sleep mode.
Tom Kurata: how fast can it come back on? If a packet comes in, does it have an auto turn on? If user intervention is required you might miss something. My solution is that if you are on the network you don't go to sleep - you can't afford to miss anything. At MAC I don't know enough - if you're logged on to the network you can't go to sleep.
Ian: when you're a network active participant you won't be able to shut down your receiver. I have to wake up for every packet to determine if it's mine. You can't turn off - the operating system people are already trying to handle sessions being broken due to sleep. The OS has to understand that - give warning upwards to allow the upper layers to handle this.
Simon: if you want power down in PHY layers the request must come from above the MAC.
Ian: if you don't, your session will get terminated by most OS's. We could provide a mechanism that will disable this - some alternate intelligence up the stack is implicated. We can provide the primitives but they must go up the stack also.
Wim: 2 approaches: In the MAC we don't kow what the session activity is, we don't know when the next packet will arrive. This knowledge is higher than the MAC. We should provide facilities so that that can be done in furture. OR, in the MAC, have some generic mechanism that allows us to power down from time to time in such a way that higher layers can (currently) still recover with our procedure. A generic solution where power can be down under control of the MAC - but ideally the best way to do it is above the MAC. That is where the knowledge is.
Bob C.: is this in interface issue? It does illustrate the value of that interface. If management comes down the side to the PIL either of these can be implemented. That is the reason for this interface model. It is a model that is very flexible. We are

getting hung up - this isue should be on a list of features provided or not by the MAC. This structure supports either.
Simon: we are getting off the point.
Francois: there is already an issue, is this a new one? 13.3 and 13.6 and 13.7 - does this match any of these?
Simon: 13.7 refers to RF power control. That means amplifier RF power. What about 13.3 and 13.6?
Francois: will enter these pros and cons to 13.3.
Simon: back to the interface - let's hone down on primitives. Are these totally isolated services access points ... MAC will be able to control parameters frame by frame. There is a set of management primitives and a set of data primitives. When it is a data request primintive, does it have information the PHY should use to set up parameters - or should it be totally independent?
Jim: the management set should be used to set parameters in a none realtime basis. Data through the data SAP causes parameters to be loaded within the PIL but doesn't contain information that may change something. Everything is set up before hand and data uses these parameters set. Nothing changes realtime in terms of power or antenna.
Simon: so the information flow over the 2 SAPs is independent. Things can only change in the PHY between frames.
Jim: parameters get loaded when data comes down - new parameters that change something. When a command to change antenna comes down, nothing happens until next data.
Bob C.: agrees with Jim, but action could have been taken on that control information. Management SAP sends data interspersed between management commands, ratio could be one to one or spase. Between a control request and confirmation, data cannot be sent on the other SAP. The PHY SAP request/confirm time is just data time.
Ian: management MAC to PHY may take longer as Bob says, e.g. channel switching. This needs to be considered. We should list the pros and cons of putting control into same frame as data - a corelation issue.
Simon: there is an issue about sending management information whenever you want. There is a synchronization issue with data being sent at the time. But if data primitives have control options attached to them it might be easier.
Jonathon: likes the latter because multiple PHYs might require parameters to arrive at same time as data. Opens up a lot more freedom for the PHY. Asynchonous control and data it puts some constraints on the PHY.
Ian: inbound direction the information is synchronized with data?
Wim: seperate but in sync.
Simon: options: two SAPs, over which primitive flow is synchronized; OR data primitive has information attached to it saying this is how to send this. Or on receive, rx data has information attached to it.
Mike: information with data is good, but could it go through the other SAP too.
Andy: could use dummy frames with no data that are actually control going through the data SAP.
Jim: management SAP could be asynchronous control information - since asynchronous, it could be sent any time, even simultaneously with data, but n confirmed for a while. There needs also to be some synchronous way of controlling - why not have a control field in the header of the data frame (please keep it short).
Andy: dummy data is the header philosophy. Some things that have to be synchronized have to go right away. The manager has to be smart enough to know how long to wait for things.
Simon: then why not control the PHY totally using dummy frames?

Bob C.: some actions that take longer than frame transmission time - those are the things that go over the management SAP.
Simon: like what?
Andy: moving the antenna. The MAC management entity knows how long that will take. He waits the appropriate time, doesn't expect response on data gone.
Jim: why have the management SAP then?
Simon: that's the question.
Bob C.: having management SAP allows management to be at several layers in the station - higher layers could do it. That can't happen if it goes with data.
Ian: whether synchronous or asynchronous to data, what configuration to use to send frames is the issue. These are asynchronous events. The time at which the primitive is updated versus the time at which it is used. When does the old value become the new value. Frames may go with old configuration before new configuration can be used. Not just what they are, but when they are adjusted versus when they are used is important here.
Simon: that suggests that the seperate SAPs is by far preferable.
Ian: adding more SAP primitives is easier - adding data header bits is bad.
Simon: we have just agreed on 2 SAPs. I can send a request down the mangment SAP, and get the confirmation back. But until you get that confirmation expect any data to go with the old configuration. If you care, you will wait.
unidentified: then what is the difference between the primitives on the local management SAP and the SMT SAP?
Jonathon: local management SAP is specific to wireless PHY. The SMT is generic.
Wim: we can't answer that question until we explore the dotted line between the MAC and the local management entity. There is a tight relationship between MAC and local management.
Simon: so tight they are probably the same.
Jonathon: this is a conceptual box.
Simon: it's a reminder that there is a management function there.
Wim: but will there be an interface there.
Andy: maybe we need to come up with what needs to be managed and it may become obvious.

Lunch break: 11:55 AM to 1:30 PM

Functional MAC/PHY Interface Requirements, IEEE 802.11-92/78, by Wim Diepstraten

This document provides inpout on the functional characteristics of the MAC/PHY Interface.

Variable PHY preamble - we may have a whole range of PHYs from 1 to 20 Mb. High speed PHYs need more acquisition time than low speed. This is could be set over the data SAP.

Simon Black: are you suggesting the MAC puts training sequence into an empty MPDU?
Wim: the sequence of events over the data SAP could be: send preamble; preamble sent; send data; data sent.

Support for ack protocol - there can be unique functions related to ack protocols. For instance a system where the receiver selects the best antenna, for the return part it could be good to use the same antenna pair utilizing the reciprocity in the channel.

Support for dynamic power control - document 92/76 deals with that. Apart from tx/rx data, a power control entity to set up the rx threshold and read out rx level is required. From the MAC you need some type of control that allows power amplifier control and threshold comparitor and the ability to read out rx level and put it into the data received. Receive level, silence level and other things related to that need (quality of service indication, signal quality, bit rate, antenna diversity information) to be tied in with the packet received.

Bob Crowder: level of resolution in power information? 16 bits? 2 bits?
Wim: haven't thought about it, probably would require a db granularity in the range 15 to 20 db. Those are subjects for more detailed study (receive level is the dynamic range of the receiver).

As part of the access protocol you issue a command to the PHY where you want the power level and thresold. Then allow a certain time to effectively measure the signal coming in against the set thresold. The MAC needs to transfer that power and threshold level before starting to access the medium.

Dynamic bit rate selection - for different stations. For a far away station with large delay spread you might be better off with a lower bit rate, closer stations can be serviced with a higher bit rate. For that you need a table somewhere in the MAC that says bit rate to be used per destination. The same type of table can be kept for power management - per destination transmit power level.

Simon: how on a per destination basis? How do you know?
Wim: in the quality of service indication you request auto-rate detection in the receiver. On a per packet basis - when the symbol rate is the same it could be done - just the number of bits per symbol. Different symbols can mean different modulation schemes, so that is hard.
Jonathon Cheah: at the MAC when you clock that in how does he know?
Wim: the PHY knows - it could be in the PHY preamble.
Jonathon: dont wan't PHY to read the preamble ... we can discuss this further later.
Bob C.: during the preamble the receiver could unambiguously determine the bit rate? So accurately that no acuracy in bit error rate is ever received? Are we trading off error rate for this feature?
Wim: it should not have to cost bit errors.
Jonathon: if it can be done, it should help bit rate because you choose what is best for you.

Power management/sleep mode - discussed earlier this morning. Still up for discussion.

Frequency selection- we discussed specifically with the example of FH as to who would do this, but there will be arguments for both PHY and MAC doing it.

Simon: what does the MAC know about frequencies? It might know about channels.
Wim: the MAC should know where that change of frequency ocurs somehow. Perhaps the frequency is irrelevent, so whatever is used the PHY could decide. But could it left it up to the MAC - it could be easier to synchronize there. Perhaps people more familiar with FH can shead light here. Tight interface needed between MAC and PHY to ensure synchronization between BSAs - that might be a reason to put control into the MAC. Roaming in a hopping environment may need more control to montor other AP's - another reason to put into a MAC or a MAC-like layer.

Tx/rx control - determines access to the network. Decides when to switch to transmit mode.

Simon: if I send a data primitive to the PHY, isn't it the PHY that does the rx/tx switching?
Wim: yes, not a management function.

Quality of service per packet - receive level, signal level, silence level, diversity information.

Intelligence level of MAC and PHY. The MAC should have control over the PHY and has enough information to make decisions necessary to control the PHY according to quality of service.

Bob C.: does the MAC on the same PHY have to have dymanic power control, bit rate and frequency selection and variable preamble length or are those a way of saying alter the media - one common command?
Wim: for preamble length for instance, that will change for different types of PHY, while bit rate selection could be universal or a per PHY option. Power control has most advantage in a single PHY. The scheme I have in mind could be mixed - stations that use or don't use power control in the same BSA. It could be an implemenation option. This could be advantagous in mutli-channel PHYs
Steve Chen: How is signal quality defined? Is it media dependent?
Wim: in the medium dependent layer it will be media dependent. It will be converted in the independence layer for the single MAC to something generic indicating the quality of the signal.
Steve: how would you come up with this?
Wim: we have this in Wavelan, but I can't describe it.
Simon: at the management interface the quality of service parameters have to be medium independent. The problem is that to be able to adjust a given PHY you may have to know specific information for that PHY. Can it really be medium independent?
Wim: that's an issue.
Jonathon: the parameter is the same - say you have 1 to 10. That range is common. If the IR PHY is better than the radio - it will still report on the relative scale.
Simon: how do I know what to do based on that?
Jonathon: pick what you want to optimize on and maintain it there.
Bob C.: the PHY has to pass along the best it could ever get, the worst it can ever be. If not how do you know what to go for?
Jonathon: if bit error count hits a thresold you set the quality of service. It betters the level if it can.
Steve: so bit error rate = signal quality.
Jonathon: it is a measure, but not necessarily what we want to do.
Wim: that would be one obvious way. Converting that to a level would be a media independent layer function. Could also be influenced by things like delay spread as well.
Simon: from the MAC point of view, if I get a frame with a good CRC, why do I care about quality of service?
Wim: optimal bit rate choice selection for instance. I don't have a pat answer for that.
Jonathon: if you get a good frame, you don't do anything. But once you have a number of bad frames and you may want to optimize according to this rate. You cannot take this action to rectify unilaterly - the other guy you are talking to has to do it also.
Wim: remember the PHY receives every message in the air. I deal only with those directed to me - the signal quality on those.

Jonathon: you have more information than you think you have. The MAC has the filter for all packets. You could be accumulating statistics on all frames received. When you receive a frame and its ok, before you reach that level you have a lot of information.
Wim: when you look at counts maintained in the PHY, they are not that interesting because you only care about those relevent for messages directed to you. So any link counts have to reside in the MAC beauce he is the only one who can maintain counts on things for him.
Simon: management primitives can go to the PHY for action, and you get information back. Implemented on a frame by frame basis. Now how do I relate information received to my packets?
Wim: there must be good synchronization between the two. When a packet is passed to the MAC you attach information to it.
Bob C.: there is a major block missing - the media watcher. You are not describing the MACs that I know.
Jonathon: that is the reason we have so much misunderstanding this week. In the management point of view - every frame received may or may not be for you.
Bob C.: in other LANs any bad frame is just junked. Now you're asking us to maintain statistics on these and on ones not for us.
Wim: no, only messages for you are important. MAC has to go through every frame, you can look at the information in every frame and gather information.
Bob C: this is the media watcher - watches and gathers this information from this.
unidentified: do you keep track of this information that is not for you? This doesn't buy you anything. You could keep the information for all the people you are talking to. But what is the point of him keeping this information - FDDI did some link quality and that's ok if the medium is the same. In a wireless media this is the same.
Steve: why is dynamic bit rate selection good?
Wim: supposing you're operating an x-meg system in a factory. A set of your far away stations experience echoes and others close don't. That far away set could change rate.
Carolyn Heide: also to maintain compatability of equipment. A low speed station must not be able to disrupt conversation of higher speed ones.
Jonathon: robustness should be the only reason for changing.
Carolyn: at the very least you must have a way to tell the slow ones to shut up.
Tom: the MAC has to do this negotiation, like faxes now. I don't see dynamic bit rate as being determined by the PHY.
Bob Rosenbaum: muli-speed conversations going on at the same time. Speed control on a packet by packet basis - so the whole network doesn't drop speed to accomdate slow stations.
Bob C.: fax point to point is different.
unidentified: current autobauds are slow - kbits. It is a connection service - to complex.
John D.: changing modulation may also be involved in changing bit rate.
Wim: yes.
Chan: the value of the effort to have adaptive rate will be minor. 3 db lift from 1/2 the bit rate will save very few messages. Changing the rate will not fix range problems.
Wim: in a noise, channel limited envronment it might buy you something.
Chan: if a large per cent of your messages are successful, the number you save by dynamically changing your rate is negligiable. There are other remedies for this - rate change may cost too much to be the pratical change.
Wim: issue - does the standard support variable bit rate.

Jonathon: this is good suggestion. In a product point of view it costs more to put 2 antenna than change rate. Another way of skinning a cat.
Wim: diversity is for fading effects. That is another issue.
Simon: its an important issue, but let's move on.

Return to presentation: PHY managment characteristics are listed on page 3.

Simon: mixed list? On request the MAC ought to be able to get some of them from the PHY. But others are on the management entity side.
Wim: maybe PHY service specification. This maybe for the PHY media independent layer.

This is not necessarily a complete list, but Wim has thrown in what he thinks applies.

Francois: issue 18.4 is the varibale data rate issue.

Security means PHY level security. What does that mean? PHY cannot reject signal from unauthorized devices. We could use an encryption algorithm - is this PHY or MAC?

Simon: we have 802.10, this is irrelevant.

End of Wim's presentation. General discussion begins.

General Discussion
Simon: when I receive a data primitive in my MAC from the PHY. Is the information part of this data only. This is easy to see in the transmit case, but not in the receive case. And if the CRC fails what do I do? Under what conditions do we change these - what algorithms?
Bob C.: write some primitives down and see how many of the things proposed can be encoded in them.
Francois: primitives - I proposed a set that may be limited. We want to discuss the parameters with that set. The primitive names that Bob has, or mine, who cares. We want to talk about what is carried by those primitives.
Simon: shall we try to define that set?
Jim: its in Francois's paper.
Simon: but given those types and the SAPs we have defined, can we identify those pirimtives?
Jonathon: not without the PHY - it would be a waste of effort. I think this should be PHY driven. They provide a set and we act on them
Francois: for a start, the ph_data must be there - no one will argure that. But what the parameters are we will argue.
Paul Congdon: agrees we should wait for PHY.
Simon: yesterday we asked them to think of media independent objects and I guess the primitives are dependent on those. Also, how easy is it to define these in a group like this. A submission based this would be better. So, lets take a break and then do more papers.
Jonathon: since the model has come so far, let's pour some quick drying cement around it. We don't want to make tearing it down easy.
Simon: how do we pour the conrete?

Motion #1: **That the 802.11 plenary vote to adopt the protocol model developed in the MAC working group.**

Moved by: Mike Bergman
Seconded by: Bob Crowder

Motion Discussion:
Paul Eastman: while the plenary vote is not binding, it's 2/3 vote would be required to change it. This is a bit of concrete.

Approved: 29 Opposed: 0 Abstain: 4 ***Motion #1 passes***

Break 3:10 PM to 3:30 PM

Performance Simulation of Transmission Sensed Protocol, IEEE 802.11-92/72, by Jonathon Cheah

The motovation for this document was an article in a May issue of Network World which said why doesn't 802.11 conern itself with PHY level and borrow MAC from 802.3 or 802.5 where chip sets are readily available. They alluded to the fact that the market is there for cable replacement. The intention here is to show that 802.3 MAC had problems in the wireless envronment. Responding to the industry demand would be the wrong thing to do. Wrote to them and there was a small article published (Carolyn Heide pipes up to say that it was a good article). Vendors and bosses want to produce a LAN from our efforts fast, an easy way out MAC is not the right approach.

If you want a CSMA type MAC then this is the wrong market - if it will work in a radio LAN it would have taken off with the products out there already. If the market is there, the product is there already. Most customers don't care about standards, they would have bought it if it was good enough. Throwing all caution to the wind - don't, we need to be cautious.

We have seen a number of good simulations - remember them, they represent a lot of good work. A good survey has found out that a network cluster should be above 50 stations. Bear this in mind in the MAC design. Simulation is easy to do - last night we saw a nice platform - but those simulations are geared towards the cable environment, even those that say they have radio. To look at a MAC protocol - what are the things to be concerned about?

Reference #2 has a lot of effort behind it, and in this submission there are numbers gleaned from that. What does it mean? A theoretical summary of what it shows - in a superficial sense, off the cuff, "can or cannot" is not absolute. Take this not as black and white.

We have been concentrating on the ISM band. While the frequency goes up the numbers can be modified to get different pictures. Hopefully someone will now have time and facilities to measure at each frequency band and create another database. But until then we have to live with what we've got. 5000 sq feet is a survey piece of data for an average wireless LAN size.

The table on page 4 compares a MAC with a central controller and one that is completely distributed. Points about that table:

no matter how good at code isolation you are, if you have no power control you can't do code isolation.

Wim Diepstraten: no power control possible?
Jonathon: no power control is possible if your dynamic range is very large.
Wim: couldn't you do power control in away that would allow code division multiplexing?

Jonathon: yes, I stand corrected.

if you have central control you must have some means to do frequency isolation. If you don't do power control the rejection would be much less than your power variation.

trying to do a cheap station. If you use a serial delay lock loop it is the cheapest way and that introduces a lock delay corelator. It will be more expensive than a traditional lock loop, in that case.

equalization delay as well in distrib. (ISI = inter symbol interfence, and it should be ICI = interchip interference)

even if you don't synchronize carrier you have to sync bits for distributed MAC. In centralized MAC its done already.

Wim: you're not comparing MACs, but a topology all traffic goes through.
Jonathon: centralized means there is a heartbeat always present. Distributed means ad-hoc non-coordinated, no timing constraints, just burst when you want to.
Wim: fully distributes alows peer to peer, and centralized only talks to a center point. If not, I don't understand the training and lock elements.

listening to a heartbeat gives you enough time and information for frequency reference and bit lock.

Wim: how do tick marks help - how do ticks from Carolyn help when I'm talking to Simon? They only help when I'm talking to Carolyn. So all messages must be going through this central point.
Jonathon: yes, you're right.
Wim: it is a centralized topology and a distributed topology.
Jonathon: could be centralized and distributed at the same time - when talking to Simon you go to the distributed topology, when talking to Carolyn you have to centralize. The point is not to credit and discredit either - it is to point out nuances of certain topologies.

There are trade offs with expense to save delay time. Also, terminal speed is missing. We discussed it a lot a long time ago. From reference #3 the speed is a lot higher now - in the original it was lower than it is in the requirement document now (10m/sec). Another little delay you have to contend with.

Paul Congdon: speed: What does this mean? What about a terminal in an airplane?
Jonathon: speed of terminal (for instance on the GM assembly line) introduces another diversity.
Wim: statement says there is a problem with collision based MAC.
Jonathon: CSMA.
Wim: what about acks so that you recognize when it didn't get through.
Jonathon: you don't know why he didn't ack - collision or interference. You should take different action and you don't know how.
Wim: this is not unique for collision based MACs.
Jonathon: but it causes the most problems for them. If centralized control I have ways to overcome it. In collision based I don't know what to do about no ack. There are areas of concern we haven't heard for quite a while, I am trying to them point out. It is on the record that some people wanted to do - you might want to think twice about these things.

Cable replacement markets - wireles RF LAN has failed. People with aspirations to go into this market should think about it and concetrate on the services cable LANs cannot provide. This is my personal opinion and its worth absolutely nothing.

Bob Rosenbaum: this paper is a combination of technical information and market suppositions. Steve Chen: yet the title is performance simulations!
Jonathon: I didn't want to do something important - I wanted to address a particular article then looking back at a simulations I saw contradictions to what I had in mind. They lacked the propogation element of the WLAN. If these characteristics are benign then there are standards we could use in 802 already. The reson we are here is to overcome the hostile environment in the channel that we have to take into consideration. Simulations based on benign channel characteristics don't serve us - I am concerned that the business people that look at what you can get for the money would be misled.
Bob R.: this paper has nothing to do with what you just said. Your paper says current result in market have failed.
Jonathon: if cable replacement had succeded we would see cables being ripped out and we don't.
Bob R.: there are not many companies doing cable replacement.
Jonathon: the industry consultants are saying that this should be for cable replacement.
Wim: you say that the simulations done so far lack sufficient propogation elements. Offline can you specify these to me so I can put those in.
Simon: to me too please. In the simulations we have been doing we are moving gradually to a more and more realistic model. So far the little twists haven't had much result - all they increase is the simulation time.

The Structure of Slotted ALOHA DAMA MAC, IEEE 802.11-92/77, by Jonathon Cheah

Submitted the protocol proposal 2.5 years ago and people said we need numbers for simulation. When you deploy that network you have a configuration file - it should contain information as per this document.

See the original submission for protocol details.

I don't have time to do simulations myself, so here are numbers you can use.

Bob Crowder: is this broadcast for control information?
Jonathon: no it is just broadcast data. But if the upper layers were so designed you could use it so.

High-performance Access Control Method for Base Station-Controlled Systems, IEEE 802.11-92/71, by Yoshihiro Takiyasu

The document annex "802.11 July 92" contains the overheads used for this presentation [sec note: each slide is referenced by title in the following paragraphs]. Some words used in the definitions for this paper may be different from the standard words used by this group.

The acronym for the MAC protocol proposed here is BLMA - Bandwidth-request Labeled-slot Multiple Access.

System Model and Communication Paths: in the system model depicted path 'c' can double the efficiency from 'a + b', but has a higher ratio of hidden terminals.

Frame and Fragment Format: the base station transmits frames with FA containing information such as "this is a new fragment" or "this is re-transmitted fragment". Base station transmits to assigned terminal id (i.e. source address) using this frame and fragragment format. The frame efficiency is about 75%.

Error Recovery Control Strategy: for using exitsing LLCs BER must be less than 10^{-8}. 10 Hz is the Dopler frequency - lower than the fragment frequency. So fragment errors should be dealt with as random errors. Roughly, using only FEC in H/W to decrease coding effort.

Error Correcting: where m=15, statisfying quality requirement by using only FEC, it is clear that an even more powerfull code must be adopted.

Optimal Fragment Length Characteristics: transmit efficiency is reduced due to fragment header overhead - ARQ lowers the efficiency in low BER conditions. So hybrid ARQ is the optimal MAC and PHY if BER at MAC and PHY is about 10^{-4}.

Hybrid ARQ Fragment Structure: fragment is retransmission unit, information is for error correction. Optimal length of information block and fragment can be employed indepdently.

Wim Diepstraten: what is the length of error correction code related to the information block?
Yoshihiro: about 500 bits information using BCH coding, 9 bits correcting code every 500 bits.
KC Chen: BCH code, so total length is 511 bits. How many errors can be corrected with that?
Yoshihiro: 1 bit correction. Therefore the code is 9 bits.

Retransmission Algorithm: if fragment response field collision or redundant re-transmission ocurs - these can be solved by carrier detecting in the response field. "Stop and wait" can achieve selective reception. MAC can decide whether to receive in accordance with the last fragment.

Communication Procedures: terminal requests, base station assigns slot; fragment received by both base station and destination. If destination can't receive the base station retransmits. If the base fails to receive also, it assigns a slot for retransmission and the source retransmits.

Access Delay Time Characteristics: these are simulation results. New request is prohibited until all re-tries are finished. The simulation assumes ideal conditions.

Conclusion: total throughput is defined from fragment utilization and frame efficiency and re-transmission loss. Fragment loss is not counted because it depends on distribution of message lengths.

KS Natarajan: why is "stop and wait" most suitable?
Yoshihro: using this re-transmission buffer control it is easy. Wireless LAN environment shorter than the transmit delay, so this runs with higher performance.
KC: ARQ policy? I calculate that with the ARQ scheme once you transmit your data, even though with BCH the bit error rate may be reduced to 10^{-5}, on broadcast the probability that you will need re-transmit (re-broadcast) is more than half. Very likely broadcast needs re-transmit a lot.
Yoshihiro: I think it is an implementation matter, but a small number of re-transmit time is expected.
KC: very likely when you braodcast multiple stations cannot receive correctly so they need to send NAK back and base station re-broadcast - this possibility is very high. So problems are introduced - multiple NAKs get back, how does it know how

many? Broadcast is required in 802.11 - there will be a problem here with the high probability of broadcast failure. How do you re-organize the active ACK back to broadcast - contention resolution is needed to solve this problem. This is not clear.
Chandos Rypinski: have you any provision for overlapping coverage between mulitple base stations?
Yoshihiro: different frequencies on different base stations. Can apply to the SS system - base station can get to large number of channels. SS may have other access methods.
Chan: will we hear other papers on SS use?
Yoshihiro: no.
Wim: the NAK scheme - when do you NAK?
Yoshihiro: when the base station misses receive, seen by error code - the DA was ok.
Wim: how do you know the address was ok?
Yoshihiro: address field maybe need a more powerfull error correction code, or redundancy.
Jim Schuessler: thank you for presenting this. A question about the battery efficiency - how much of the frame time does the receiver need to be active if there is no data for that station.
Yoshihiro: all the time.
Chan: interested to know if this is a built system or just a proposal.
Yoshihiro: not a real system but the prototype exists.
Chan: rate?
Yoshihiro: standard in progress now in Japan - ISM 2.4 GHz expected next year. There are 2 proposals: multi-channel non-SS; transmit rate of 4 Mb QPS modulation multipath fading.
Ian Crayford: concerned by the fact that you NAK broasdcast. If you take 100 nodes, broadcast, and you are likely to miss one, if you then re-broadcast, another one gets missed ... etc. You can't ACK or NAK broadcasts.
Yoshihiro: agrees. It is difficult. It is a big problem.
Steve Chen: ACK of broadcast - most of the higher layers send broadcast several times so you don't need to ack.
unidentified: ACK or NAK means ACK from receiver to transmiter or terminal to base station; or all three?
Yoshihiro: base station transmits ACK or NAK
unidentified: re-transmission not from source generally?
Yoshihiro: if base station didn't receive the fragment, then it assigns a new slot to the station and the station re-transmits. Base station ACKs to source.
KS Natarajan: preferential retry?
Yoshihiro: new request. It is prohibited until old requests have been finished. New request is prohibitted due to preferential retry.

Simon, as chairman, says that as we have another joint MAC/PHY meeting tomorrow, before we break for today, let's go back to the model.

The first flaw? If you consider the transmit case at the data SAP: send a packet of data to go; send some management information independent of that data that was sent on the data SAP. But when you receive information you need it to be associated with a data packet - to get the station address associated with that information. So you're going to get a primitive out of the data SAP (up) that has to have data and associated with that is some quality of service information. The quality comes up the management SAP. Simon can't see how that synchonization can be achieved. We are more likely going to want to put the information into the receive data primitive.

Bob Rosenbaum: why is that synchronization important?

Simon: becaue you need to associate quality of service with a given station.
Francois: you are now using one SAP? Out of the PIL receive a primitive ph_data_ind followed by data with address and CRC in it, followed by quality of service. This is then a SAP which overlaps both the MAC and the management block (in the MAC). The quality of service goes to the management, the data to the MAC. If the MAC sees bad CRC, it passes error to the management. If no CRC error, check address. If its ours, process it. The management can send back to the MAC what needs to be done, and the management can pass to the SMT statistics counts. One SAP is needed to handle this synchronization of control and data.
unidentified: what is the problem with the original model?
Simon: two primitives that have to be synchronized with each other. Maybe it's not a problem.
Ian Crayford: why restrict to two primitivess - Ethernet MACs provide as primitives carrier, collision detect, data. They are not associated in any way. While receiving data you monitor the other.
Simon: through one SAP.
Ian: why can't we report multiple events in a fashion across that interface? Why not encode quality of service.
Hiroshi: the PHY is dumb. The MAC and management can be smart independently. We want to reduce cost.
Ian: do that by not encapsulating data with information. If the PHY layer must encapsulate these the PHY gets expensive.
Francois: encapsulation - this is implementation. We are talking logical here.
Simon: thinking about primitives made us think about synchronizing this information. If you need this information on a packet by packet basis, two SAPs doesn't work.
Ian: but on transmit two is better because it is update versus use time. Passing stuff with data is a major problem for changes later.
Mike Bergman: encapsulation need not be expensive in the 90's.
Simon: this is a ***logical concept*** - where service primitives flow between layers. Well, just bear it in mind.

Meeting adjourned: 5:30 PM.

Thursday AM, July 8 1992

Meeting called to order at 10:45 AM, Dave Babgy back in the chair. Carolyn Heide secretary. Dave apologizes for his absence, due to meeting conflict, and expesses much thanks Simon Black for filling in.

MAC/PHY interface has been discussed but issues haven't been closed with the pros and cons being recorded. Jim Schuessler has an action item to respond to issue 15.3.

Briefly, on the subject of the e-mail reflectors, if you have problems fax Dave Bagby.

We also said we would try to address time-bounded services. We didn't have any papers? Does anyone have anything to say, or did anyone do anything and bring it?

Chandos Rypinski: points out his MAC proposals address that.
Jonathon Cheah: ditto.
Dave Bagby: this subject has other than MAC impact.
Jim Schuessler: had hoped we would be able to do more on this subject. National Semiconductor is supportive of the effort and is working on it. Hopefully they will contribute more later.

Simon Black: time-bounded services are interesting. The longer we leave it the harder it will be to tackle. They have massive impact - the longer we delay defining them the more difficult it will be to provide for them.
Dave B.: hears good intentions and interest but nothing to attach to it. Please poeple, make submissions - not specifying the bit alignment, but what does it mean - what does it mean to the MAC/PHY interface for instance. Abstract support issues.

Last meeting we set two topics - we did well with one and not with the other. Maybe they were too much for one session. Should it be more than one topic per session, or should we go for only one? What was productive this session?

Francois Simon: 1 or 2 should be enough. It depends on the subjects - the MAC/PHY interface was a big subject and there is still a lot to be done. You can't fix one or two subjects per session, it depends on what they are.
Dave Bagby: attacking the interface was so that the subgroups could work in parallel. Is there a lot more joint work to done?
Wim Diepstraten: depends on the subjects - frequent liaisons are still needed.
Dave B.: joint time was productive? (the general concensus is yes)
Carolyn Heide: a lot more time is still needed for the MAC/PHY interface.
Dave B.: let's not just stay on that one alone. Let's move on to another major subject, but still keep the others.
Jonathon Cheah: there is some merit to continuing on the primitives on the MAC/PHY interface. Also we have a MAC model which is well accepted. The next test of if the primitive concept works is if you want to pass that primitive you have to determine the primitive structure. Very important. Number of bits per parameter is very important also e.g. at least 6 bits per signal level.
Dave B.: for implementation yes, but now just knowing what primitives - what types of information - we need now is enough. If we don't talk about other subjects we will miss things on that interface. If we go to the bit level and define in detail we may have to change drastically when we move to other subjects such as time-bounded services.
Jonathon: we have a framework, but how to make that framework more solid is a tough thing. Let's build from what we already have. This is a chicken and egg situation.
Simon Black: agrees with Dave. At this stage we have gone too far a couple of times into bits and bytes. Where the lines are and what the objects are on either side is sufficient for now. We don't have enough information to get into the bits yet.

Next meeting:

1. MAC/PHY interface:

(a) identify objects passed across (think about is there anything not said yet, and why);

(b) time-bounded service implications (not how to do it, but things that affect the MAC/PHY interface, what and why)

(c) distribution system services (without that all we have is a single BSS system. This is related to both the above. What needs to happen, what kind of objects, what kind of services. Avoid bias to your own implentations).

This is already more than we will be able to handle. This will drive our agenda next time. As I said last time, we will try to accomodate non-item papers - but priority will be given to submissions on these items. We need to focus, we can't deal with all subjects at all meetings. Non-item papers may need to be handled at another meeting later.

How much time, as a percentage, on (a)? There are suggestions of 25%, 60%, and 30%.

Simon reminds us that we should make sure issues addressed are specifically addressed. Dave says we may break into groups of 2 or 3 to list the pros and cons of specific issues. Contentious issues will become obvious. This kind of working method is appropriate for an interim meeting.

How much time on (b)? The only suggestion offered is 25%

How much time on (c)? No suggestions offered, but Bob Crowder suggests we should spend more time on (b) than on (c).

Chandos Rypinski suggests that we leave unallocated discussion time too. Wim Diepstraten points out that the amount of time spent where depends on the contributions. Dave says that those people who expressed and opinion on % should submit papers which would reflect that.

Get e-mail access if you can. Compuserve is one easy way. The way the reflector works - you send a message to it and that message goes to everyone. We all have trouble doing any work outside of meetings - but we won't get anything done if we don't. This is a discussion outside of meetings tool and using it doesn't take up a lot of time. MCI has a policy that says no mail forwarded to MCI that started from MCI - don't know any way aorund it (yet). Suggestion from the floor- somehow re-send mail with the reflector's address so that it no longer has originated with MCI. Dave will investigate trying to do that. It's just an alias now, would have to get it more active. He will see what he can do. In silicon valley there are places from which you can get Internet service for something like $15 a month, so its not very expensive.

An agenda with these subjects decided above will be issued. Will set the balance of time according to paper submissions. Much thanks to the people who submit papers because that's the hard part.

Anything else? no, what else can we address in the time remaining? Wim Diepstraten suggest that we may want to change the model before it is submitted to the plenary. Ok, we agree to discuss that some more.

> Simon Black: the model was useful but we got bogged down trying to classify the information that flows across the MAC/PHY interface. We don't have enough information to propose seperate SAPs. The local management box is very important in itself, but does it need its own SAP?
> Bob Crowder: you are suggesting that there may be primitives destined for it but not through a different SAP?
> Simon: needs to think about it more - not saying no, just that we don't know yet. What the structure of the primitives is is unknown now. Not convinced that we know enough to organize the MAC/PHY interface yet.
> Bob C.: is uncomfortable not with taking out a SAP but with removing the concept that there are 2 different types of primitives. We don't want to loose the idea that there are different time domaines on the responses from primitives sent to the PHY with respect to when it will get back a confirm.
> Wim Diepstraten: currently we have examples on the table which would benefit from a single SAP. We have examples where a tight synchronization between the two is required and can best be dealt with by one SAP. Both the information sources can be there - MAC and management - but accessed through the one SAP.
> Dave B.: There's data and there's control. I am thinking functionally and I'm afraid you're think implementation. One path or two paths doesn't make any sense here. If I'm messed up here I'm sorry.

Simon: the point, which is from a discussion we had that raised the issue 1 flow or 2 flows, was if an MSDU is received there may be information purely related to that MSDU. The other thing is there is information that is independent of frames. The management functions in the MAC were because some of them depend on which station is being addressed.
Wim: an extra aspect is synchronization needed between management entities and data entities. As Simon explained there might be management information attached to a receive frame and it is important that that relationship is maintained. That's why one SAP.
Dave B.: synchronization issues?
Wim: yes.
Francois Simon: local management is important.
Andy Luque: one or two spigots on the logical interface - make a list of what goes across here and then look to see what is the needed absolutely. It's a religous thing now.
Dave B.: would you bring a contribution that does that?
Paul Congdon: submission 92/61a from Nathan Silberman has a lot of that already.
Paul Eastman: other standards may have similarities we should check.
Francois: there are not likely to be many existing examples.
Bob C.: two ports are used in the field bus example presented. You haven't had chance to read it yet, but please do. It is defined there. I said during that presentation that it is possible and normal to pass information in the data stream. Keep clear in your mind that there is also strickly management information. Jonathon Cheah proposed a generic MAC, and a PHY specific part of the MAC - you might see the two SAPs fit into that concept. I am in favour of 2 SAPs - we have made good progress and we shouldn't back up.
Mike Bergman: agrees with Bob. As to what Andy said - make a list - we have a list of 3: data, control/management tied to data; control/management independent of data. Keep 2 SAPs in the model, but don't forget that control can go through the data SAP. The first 2 go through the data SAP and the latter through the management SAP.
Dave B.: control with data needs to be issue.
John Deane: is this a functional or an implementation issue?
Dave B.: functional unless the interface is exposed.
Bob C.: yesterday we said that DTE/DCE interface is where exposure ocurs. Not the MAC/PHY interface. We discussed that.
Wim: we have a new concept called local management. We are discussing SAP use. OSI management has a seperate SAP - the field bus specification has OSI SAP management seperate.
Andy: SAP for OSI management - one aspect is communication with peer entities in which case the data stream SAP is used with an MPDU. A 3rd SAP for OSI management - use the thin line between OSI management and local management - use that side as opposed to a 3rd SAP for the MAC box.
Simon: the electrical interface is wholely irrelevant to our discussion here. A MAC/PHY interface is being disucssed here.
Dave B.: when I first saw this diagram I thought I didn't understand the DTE/DCE interface. The MAC/PHY functionality is being discussed here. Electrical plug interface is traditionally how to connect a cable to the media. In our case the medium is always connected. You cannot define the equivalent plug. There's another level that is a plug that connects 2 black boxes - one MAC and one PHY. This is not being addressed here - that is a dangerous discussion. It doesn't fit into a this discussion.

Bob C.: we disucssed this. Some of us see a need for this DTE/DCE discussion. A limited number of pins can accomplish this and it is the only way you will ever make one MAC and many PHYs work.
Jonathon Cheah: subscribes to the DTE/DCE concept. You assume it means we want to have an exposed connector, and this is not necessarily so. It is a formalized connection that may not be so. Unexposed connection is up to indivdual implementers. We want the interface to form a formal seperation between PHY depedent and independent.
Paul E.: even if proposed and not exposed it is usefull to have this interface defined so that with a proprietary harness you could attach something. This is needed for conformance testing.
Dave B.: two issues: (1) is there DTE/DCE, and do we define such a thing?
Bob C.: we voted to recommend that model to the plenary.
Dave B.: we have a procedure for issues and anyone can open one if the formal procedure has been followed. The other issue is (2) will that interface be exposed?
Jonathon: the issue list update information can be read from the meeting. Both sides have been statisfied.
Bob C.: we followed the procedure. We discussed issues. We voted unanimously to take the model to plenary.
Simon: reminds us that we have a procedure that says anything we come up with and think we've fixed, we open an issue and let it sit for one meeting. We can't adopt this at this meeting as fixed. It has to remain open. Some people have some concerns with the model - electrical interface is only one, there are others. We can tell the plenary it currently has a lot of support, but we cant close it.
Bob C.: no technical decision can be taken in one meeting?
Simon: yes. A lot of people have said they want to think more.
Jonathon: to make progress it is important to make one step at a time. The intention when the model was done (my intention anyway) was that if you don't put up something and shot at it you wouldn't get far. That was my intention when we made a concrete vote to adopt this model as a baseline. This is the starting point and we got a unanimous vote (which is most unusual) - to throw that away because of procedures - let's put the good of the committe first. Let's not argue about what we have decided.
Paul E: this has been in dicussion for well into 3 years now - this model. It's wrong to say you shouldn't make any decision in any meeting. Perhaps make one, then confirm it next meeting. Also, if you chose to re-confirm at the next meeting it is not a yes or no, it is yes or move to a proposed alternate.
Dave B.: what if one came to our next meeting and wanted to see this picture and read the description of what this picture meant?
Bob C.: I thought Simon Black volunteered to write it up.
Jonathon: the issues editor and the secretary will have captured the essence of the discussion. If not, then we need a better method of working. We had a very productive discussion and it would be tragic if that was lost.
Dave B.: 2 answers - in the minutes and in the issues log?
Francois: there were no issues raised. We have not resolved any issues at this point.
Jonathon: but the information is there to satisfy the procedure. We have the information in the minutes and could pigeonhole it correctly.
Bob C.: I would do it if Simon doesn't want to.
Simon: I would like open an issue - is this the MAC protocol model we want adopt?
Dave B.: procedural problem - the way to open and close. You can't close an issue that hasn't been opened. The concesus on this model involves several issues which has not been phrased. There is nothing written down specifically - I would have to ferret out unspecific information. This has caused problems in the past. I personally

suspect that everyone has different ideas about what this model says. There will be arguements about what we adopted. Spending time formally specifying this would be worth the effort.
Chan: my understanding of our vote is that there was a decision of the MAC committee that there was a further step to resolve at the plenary.
Mike: treat the model we have as not a formal completed standard. There is going to be disagreement about issues concerning this. But as a thing to stand on it is good - disagreements are fine. All we are saying is this is the starting point. Disagreement is normal. Until someone has writen it all down the discussion is premature.
Bob C.: to back off on a decision that made progress would be damaging. Decouple the issue of the picture adopted from the description of the picture. I would be willing to work with others to put text to the picture and clearly it would not be attached to the motion.
Jonathon: if a motion is passed by the MAC committee it should be carried out unless there is another motion to cancel it. I would like to pour conctrete - the issues list is for these questions. If we close the issues the picture gets changed. We have to keep the step we have reached and not back off.

Motion #2	**To open an issue regarding the protocol model adopted in the MAC group and not to take ~~a vote in the 802.11 pleanry~~ this afternoons vote in the pleanry and move adoption of the MAC model to an issue.**
Moved by:	Simon Black
Seconded by:	Paul Eastman

Motion Discussion:
Dave Bagby: the motion and vote were made when I was not here, and it conflicts with issues procedure. I wasn't here then and cannot judge whether if I would have ruled that motion out of order. Now there is motion to undo the previous motion. This is difficult to handle.
Bob Crowder: calls the question, seconder Dan Lewis (22, 0, 2)

Approved: 8 Opposed:13 Abstain: - ***Motion #2 - fails***

Bob C.: offers to write text to go with the model with a small group of volunteers.
unidentified: opinions were expressed - no issues were opened. There was no procedure violation - there were no issues to be recorded or closed.
Jonathon: you can still open issues after we adopt this. Present this to plenary so the PHY subgroup can look at it. We have been proceding correctly in the last few days. We don't want to recant what we have already agreed upon.
Bob C.: can a group be made to do this?
Dave Babgy: anyone who wants to work with Bob Crowder on text to go with the model, let him know.

Meeting adjourned at 12:25 PM.