Joint MAC/PHY Subgroups Meeting
Tues. PM, July 7, 1992

Meeting called to order at 1:45 PM chaired by Simon Black in concert with Larry van der Jagt. Secretary Carolyn Heide.

Agenda:
- (1) Presentation of Input Papers (1 hour 30 minutes)
  - 92/69, Francois Simon (20 minutes)
  - 92/73, Bob Crowder (40 minutes)
  - 92/78, Wim Diepstraten (15 minutes)
  - 92/85, Jim Schuessler (15 minutes)
- Break (15 minutes)
- (2) Identify the issues that most affect the interface (30 minutes)
- (3) Time Limited Open Discussion (1 hour 30 minutes)
  - - allow each group to understand views of other group
  - - obtain alternative positions for the issues.

Goals:
- Identify the issues that most affect the interface
- Obtain alternative positions on the issues

Agenda Item (1) Presentation of Input Papers

MAC/PHY/Management Service Interfaces IEEE 802.11-92/69, by Francois Simon

This paper is proposing just a model. It doesn't propose MAC/PHY or Management design, but a language which may be used when we go farther to specify specific requirements.

A series of primitives which enable, in an abstract way, the communication between layers is described. Primitives represent the logical exchange of information between the layers, based on the ISO seven layer model.

Discussion:
Jonathon Cheah: assuming get, set and event - the first two are bi-directional, the latter is uni-directional?
Francois: event is unsolicited and the others are request/confirmation types - an exchange of information.
Jonathon: the MAC to higher layer - does that cover most of other LAN and other device interfaces - our MAC at this point may talk to other higher layers not necessarily within the confines of 802.11. The set must then cover any devices we may attach to. Is that three sets of lines [sec note: in the diagram on page 4] enough to cover most of the lines we would encounter in the future?
Francois: that's the way I see it, yes. Activate, deactivate, data. Management interface to higher levels is not defined.
Jonathon: what about stand by, or not ready? Like a flow control indication - this is not covered by these.
Larry van der Jagt: flow control is handled at a higher layer. Usually at the session level.
Francois: or at the LLC.
Orest Storoshchuk: what if you attach to a Novell or something?

Francois: these might not be sufficient but those other things have not been defined yet either. When identified they could be placed in. This is just the ISO general model used in CCITT.

Larry: in previous MAC/PHY interface discussions we have identified things that are not data - power, antenna selection, etc. Should that come through the management path? I don't believe that because the MAC state machines will set these, not station management. Communication enabling things will be MAC functions. Another primitive to handle non-data primitives between the MAC and PHY is required.

Jonathon: the MAC could pass state machine parameters to management and that could pass them down to PHY.

Larry: CCITT does not have information that needs to be passed on a packet by packet basis. Management is usually a S/W function.

Francois: management is also used for managing the network.

Larry: that is another layer altogether - network management software. We have no data information that is required to pass across the interface.

Francois: that hasn't be defined yet.

unidentified: agrees, there may be information such as sleep mode that needs to pass down not through management. Management should allow use of different protocols. Unless the management has hooks into the higher layers, how can it pass information?

Larry: our case is unique because pseudo-management data for us can be as important as data.

Jonathon: if there was another line between the managed objects of MAC and PHY over which primitives could pass that could be handled.

Orest: traditionally the management items go all the way up the side. Now the MAC needs to manage the PHY on a dynamic basis.

Larry: there are more things in the symbol passed to the PHY, more than just data. Within this primitive structure another primitive is needed for things other than data. Or the data primitive has more in it than data.

Jonathon: in my opinion that is under the management blocks. Another block is needed.

Larry: no, managed objects belong in the management blocks, these go right into the MAC state machine on a packet by packet basis.

Francois: I propose that we have another primitive 'ph_information' that will have types as defined here.

Tim Kwok: does the management have signaling information?

Francois: this is a generic model. Signaling has the c-plane and u-plane (control and user) there is a MAC and PHY on each. In this model there are not two planes.

ISA-dS50.02 Partial Copy, Chapters 11 & 13 Omitted, IEEE 802.11-92/73, submitted by Bob Crowder

This document is the field bus standard. The PHY layer is far along and is becoming a draft international specification. It has a 4 layer architecture. Four PHY layers are defined - baseband wire media - 31 kbit/s to 1 Mbit/s (2 of them) and 2.5 Mbit/s. There is a plan for fiber optic and radio media.

Bob is presenting this standard because it has a MAC/PHY interface that he believes satisfies 802.11 needs.

The MAC and LLC are combined into the data link layer.

MAC/PHY interface is not a physical interface it allows passing of octets down using a PHY data primitive. Frames are passed down to MAC, octets presented on the

DTE/DCE interface. Station management interface to the PHY is above the DTE/DCE interface. Management commands can be carried on the same pins that carry the data down into the medium dependent sublayer. This layer has to do with signaling characteristics. This is where we would find the issues of frequency hopping, spread spectrum, infrared - below the DTE/DCE interface. This concept allows you to build one MAC and connect to different PHYs. Below that specific PHY logic is the media attachment units (optical receivers and transmitters, Manchester encoders, etc). At this point we have a second physical interface. If exposed it must conform, but is optional. This is a very simple interface that allows you to do low cost implementation and forces media dependent stuff into PHY layer.

Larry van der Jagt: what's in octets and what's in bits?
Bob: PHY is octets, below that is bits.

Data request; data indication; data confirm - these are the primitives at the PHY interface.

Page 9 has primitives which very much match what Francois Simon was just describing in the last presentation. Also has a PHY characteristics primitive which tells the MAC what it is dealing with. It is only through that mechanism that you have a truly independent DTE/DCE interface.

Section 5.1.3 describes exactly what happens for transmitting of a frame. MAC/PHY flow control is controlled by the PHY confirm - the MAC cannot send again until it gets PHY confirm.

Section 5.1.4 describes how noise and error conditions can be determined by reception or non-reception of these primitives.

The side station management to PHY interface - certain management services are optional. It is mandatory to be able to reset the PHY. SETVALUE and GETVALUE are optional. SET gets a confirm, GET reads parameters from the PHY. EVENT reports PHY layer events - a glitch, an unsolicited occurrence. SETVALUE says that, within the range of the wire medium, the parameter set on page 11 is a sufficient set of parameters and values for the MAU. Maybe it has to be extended for a radio MAU. GETVALUE has no confirm, it is immediately executed. EVENT is primarily used for fault reporting.

Jonathon Cheah: an event indicates an error, so doesn't that have to be passed up in the stack?
Bob: note that the management primitive is on the side. It is really wherever the implementer wants to put it.

The DTE/DCE interface is a potentially exposed interface. There are 3 levels at which that interface is defined - services, messages, signals. There are no optional services here - there are some optional messages. The services are defined on page 13. The message transmission and reporting services do SETVALUE and GETVALUE - the same services do data and management. Loopback is done by use of the message transmission service. Unusual thing concerning the DTE/DCE interface - redundant media is very important so a real-time provision is made for redundant media. With one set of pins you could control multiple medium instances. This may be very similar to radio.

Jonathon: how would the system decide which media to switch to?
Bob: in field bus it is decided by the modem that first clocks onto the preamble.
Jonathon: all media are transmitting the same information at the same time and the first one that signals valid data is the one that gets through?

Bob: yes, if they were all transmitting.

PHY signals at DTE/DCE interface are on page 14.

The MAC/PHY interface is not exposed. The 2 interfaces that are exposed are optionally exposed.

Alternates available to 802.11 based on field bus:

1. look at this as a model and try to learn from it;

2. think about this as a prototype - an initial starting point and a real interface, and start to write issues against this interface. This is an ISA standard so we will be referencing this interface with changes;

Personally Bob recommends the latter. Many discussions here in 802.11 have leaned this way. It took a year and a half to do this standard after the PHYs were defined, let's benefit from that effort.

unidentified: advantages of this model over, say the 802.3 model. why is it better?
Bob: doesn't think the 802.3 model is that clear.
Larry: a lot of the services here aren't in that model. These boxes have more capability than the 802.3 boxes.
unidentified: all the services provided here are in 802.3?
Bob: the SQE set is a big problem in 802.3. This doesn't have that problem, it is much richer.

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

Intention of this document is to facilitate a discussion about the MAC/PHY interface functionality. Not completely abstract - more like the functional requirements of that interface.

Dynamic MAC/PHY management interface where the MAC has full control over the PHY directly.

One MAC over a range of PHYs. Whether there are some PHY dependent functions in MAC - we may run into this and will have to determine how to handle it. Shot for an interface on which these may exist, but may be parameterized differently depending on the PHY.

Functional characteristics to be considered: threshold and power setting on a per packet basis; information from the PHY about the interference level on a per packet basis; there might be different bit rates per destination; who does the frequency selection in a frequency hop situation?

Intelligence distribution - the MAC should be more intelligent. PHY has to be more of a service provider.

Accommodating a Range of Intelligence in the Physical Layer, IEEE 802.11-92/85, by Jim Schuessler

The paper has not been handed out yet. This paper resulted from the only e-mail exchange that came out of the reflector. Bob Crowder replied to a question of Dave Bagby's and that started Jim's idea. In the previous papers we have seen a lot of agreement on the fact that the type of entities are control frames or packets at the media independent interface.

Start: with the assumption of one MAC and multiple PHYs.

Goal: while Francois takes a top down logical service access point approach, others take a bottom up real approach, efficient performance driven. With respect to the PHY independent layer - Jim agrees with the field bus approach. This layer helps accommodate multiple PHYs. One MAC/PHY interface - but media dependent interface somewhere. Whose responsibilities are these interfaces? The PHY group. PHY is the place where the knowledge to do this is. We all have the task of the logical independent interface between the MAC and PHY.

Problems: limitations at the media dependent interface may preclude implementation at the independent interface. But a valid reply may be "I don't know".

John McKown: you don't explicitly show - there's a fork function at the media independent PHY.
Bob Crowder: the media dependent interface is shown as a block itself.
Larry van der Jagt: the PHY subgroup has a PHY layer template document which discusses a physical layer convergence protocol and does this media independent to dependent mapping. The PHY group has already said that there needs to be an independent to dependent interface convergence. The definition of that protocol belongs to the PHY group.
Francois Simon: requests from media independent layer to the media dependent with no response - it is perfectly normal to have a dummy response.
Jonathon Cheah: in the diagram on page 3 - a small box could go into the MAC that said "PHY independent" and the four boxes be changed to one that says "PHY dependent".
Jim: these interfaces may or may not be exposed. For economic reasons none may be exposed.
Larry: looked at FDDI, 802.3, 802.4 and 802.6 models when researching the PHY template. 802.6 most closely matches what you say here and what the PHY template proposed.
Jim: to Francois - the last two points on page 3, do they map into issues? My answers are here and can address those issues.
Simon Black: they are issues 12.3 and 12.4.
Dave Leeson: could this be implemented in either H/W or S/W? I can imagine a MAC that could ask questions and get answers, or one that never asks.
Jim: yes, the standard could allow either and the implementer could solve. There could be one default PHY that was dirt cheap and couldn't accommodate anything.

Break at 3:15 to 3:30 - chairman says during break think about what the papers said with respect to issue list conclusions and arguments.

Back from the break, and Vic has good news and bad news to announce. Bad news - November meeting announcement has a mistake - registration fee of $150 should have been $200. Good news - 802.5 standard is available at registration to select groups at select times.

Agenda Item (2) Identify the issues that most affect the interface

Obvious issues Simon comes up with: 12.1, 12.2, 12.3, 12.4, 12.8, and 13.4.

Francois Simon: what about the MAC/PHY to management interface? Are we considering this too?

Simon: 13.4 is the only one I could find. Yes, I think the management interface is relevant.
Francois: then add 13.5.

Simon states that in the e-mail discussion referred to by Jim Schuessler in his presentation, Dave Bagby raised two issues (12.3 and 12.4) that should be addressed. The extreme choices are set the PHY and tell the MAC; or have MAC order the PHY and have PHY respond.

Discussion of Issue 12.3 - "What is the intelligence at the MAC/PHY interface?"

Jonathon Cheah: single MAC and multiple PHY has been decided in Functional Requirements. Hard to put in MAC then, because MAC would have to know about all PHYs. So outside MAC. Then there needs to in a PHY media independent layer.
Carolyn Heide: not a yes and no question. How can we phrase pros and cons? Could we re-phrase the question?
Jonathon: paraphrase: intelligence out of the MAC; intelligence in the independent PHY layer.
Nathan Silberman: question is what decisions are made where.
Jim Schuessler: we don't need to decide 'what is decided where' but the mechanism to transfer that information on the interface. The MAC needs to ask and the PHY needs to respond - define a rich protocol between the MAC and PHY independent layer with which to control the interface.
Jonathon: disagrees - makes the layer separation vague. A single MAC wants to say "I have a packet, send it". The layer which controls the knowledge of the physical attributes needs to reside in the PHY. The MAC doesn't care how you do it, just that you did it.
Jim: agrees if you change these things over to the management interface. Maybe the appropriate place for these things is over the management interface.
John McKown: should 12.2 be settled first [sec note: issue 12.2 reads "What interfaces are exposed: MAC/PHY? DSS? DSM?"] If the interface is not exposed, does 12.3 make sense?
Larry: if it's not exposed, do whatever you want.
John: if the answer the issues12.2 is no, then we should not answer 12.3.
Francois: there could be a MAC/PHY interface in the logical, model sense without exposure.
Simon: there has to be a split of functionality, but also just a way to talk about the interface even if it's not exposed.
Wim Diepstraten: for modeling reasons to make a standard you need a functional, logical separation between a MAC and PHY. How else can we say a MAC will support different PHYs - we need this if the interface is not exposed also.

Here we break in the discussion to let Larry van der Jagt present the PHY Template as it is relevant and not all of us have been exposed to it.

Explanation of the PHY Layer Template Document, IEEE 802.11-92/4, by Larry van der Jagt

This is an architectural modeling which each PHY layer proposal will use. This document is a PHY standard that needs to be filled in. The document details 2 sublayers - see the first page of 92/4 - physical layer convergence entity, PLCE, (to MAC); and the medium layer dependent sublayer, MDS (to media). This is modeled after 802.6.

The PLCE accepts symbols from the PHY Independent interface (MAC-PHY or DTE/DCE) above the PLCE, and maps them onto symbols used by the MDS. The MDS converts these symbols to mechanical and electric requirements of the medium. Below the MDS is a mechanical and electrical interface to the media entity, ME, which is the media itself.

This maps nicely into the model which Jim Schuessler presented, resulting in the following model diagram:

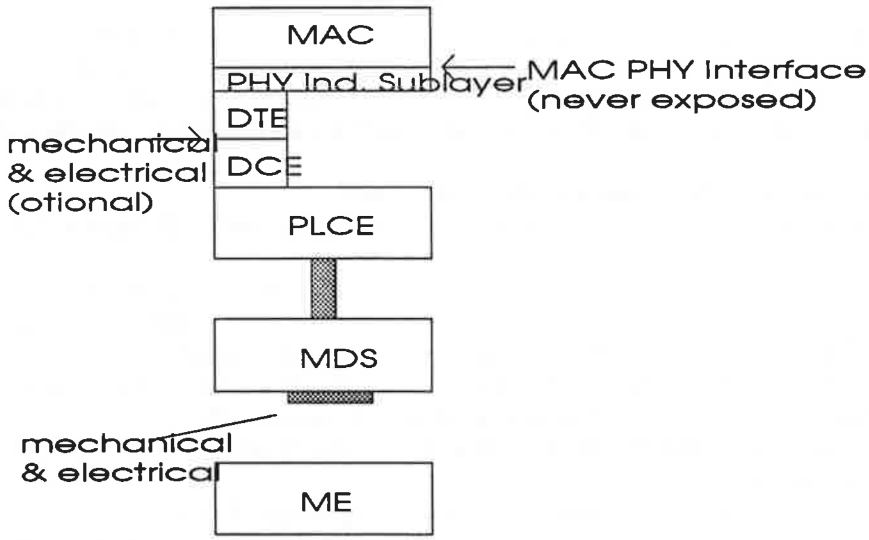

The questions: Is the top interface a MAC/PHY interface or a DTE/DCE interface which resides below the MAC/PHY interface? Is there a connector between the PLCE and MDS?

Francois Simon: what is a DTE/DCE interface?
Larry: DTE to DCE interface is basically a definition of a cable and what occurs over that cable.
Francois: so you see 3 service interfaces?
Larry: this is all within the PHY. The PLCE, MDS and ME. DTE/DCE does not imply exposed - if there is somewhere where you could plug a modem it would be there. If you have no intention of plugging in, it's logical only.
Bob Crowder: this is not accurate for field bus because management interface comes in above DTE/DCE. There is a functional advantage to that model and it's not shown that is station management coming above the DTE/DCE.
Larry: this is an argument why DTE/DCE should be above PLCE instead of MAC/PHY interface.
Paul Eastman: clarify - if DTE/DCE interface is exposed it would be a recommendation not a mandatory thing? E.G. a port on a laptop where you could plug a radio or IR PHY so you can quickly switch. You may want to buy each one from different companies so it needs to be defined. Not mandatory, so if you want to build a proprietary plug, go ahead.
Larry: there will be arguments about mandatory - 802.4 had a lot of argument.
Jim Schuessler: if DTE/DCE interface is exposed is it common among different PHYs - Paul has addressed this and it should be.
Larry: whatever we call it - that's where you plug different PHYs into the MAC.

Dave Leeson: imagine a Dick Tracey radio computer - you will not have an exposed connector. So there may be times when exposed is wanted and times when it is not. So "when" it is exposed it should be defined.
Larry: no one ever said they had to be exposed. But if exposed, does it have to be as defined or do we just recommend it?
Jonathon Cheah: being in the consumer product line, if we have DTE/DCE interface it should be optional. If exposed ok, but not mandatory to make it exposed. Since not mandatory it is also necessary to add that if an exposed DTE/DCE interface exists then that connector should be conformant with the standard. That connector allows all companies supplying PHYs to build interoperable equipment.
Chandos Rypinski: supports the plugable interface as a practical matter - PCMCIA and SCSI are the only two candidates. How much S/W on one side or the other is uncertain. For a radio on a laptop, the plugable interface will be one that is there already on the other side of the MAC. So there won't be any exposed interface other than that.
Simon Black: if we have this, how complex will it get?
Larry: keep it as simple as possible physically and mechanically. Basically the four primitives Francois described in as few wires as possible.
Simon: majority of the choices in PHY, not MAC? Simple MPDUs to PHY?
Larry: no, PHY feeds information to the MAC and MAC decides. MAC send symbols to PHY and PHY maps to real world (pick antenna, choose power, etc.). The PHY comes back and tells MAC relevant things in the symbol reported back. Make the physical transport of the information with as few lines as possible.
Chan: Larry commented that the MAC will be below the PCMCIA connection - that's not a certainty. 486 laptops ...
Larry: then you say we need a SCSI-to-MAC interface in the MAC?
Chan: if it happens it won't be because I like it. It will be because it is demanded.
Larry: there may be this DTE/DCE interface above the PLCE and the MAC above that. But if not exposed, that DTE/DCE part may not be there. Would the MAC/PHY interface be the part that might be exposed or the DTE/DCE part? Bob says that with the DTE/DCE you can run station management above.
Mike Bergman: isn't the MAC/PHY interface a logical interface not a physical one that can be exposed? In terms of intelligence we need to come up with common functions that span all PHYs. Any mechanical interface specified will be below the MAC/PHY interface.
unidentified: PCMCIA has specified interfaces that exclude MAC/PHY interface being internal - they must be exposed. Insertion loss, impedance mismatch and cross talk will be the killer issues on the MAC/PHY interface. We cannot afford loss in terms of data integrity, so this interface should be logical only.
Jonathon: if we break out DTE/DCE interface (and there are a lot of advantages to doing so) then we are talking about 2 interfaces - a logical MAC/PHY above the DTE/DCE and the thin layer that is the PHY independent layer - the DTE/DCE comes in there and the PHY dependent layer comes in below that. Two interfaces, but only one is exposed.
Larry: the MAC/PHY interface would never be exposed under any circumstances.
Nathan Silberman: two interfaces but only one thing, physically and logically represented. If there is another interface there needs to be another definition of that interface. Apply the 'kiss' principle and have only one interface - the MAC/PHY.
Jonathon: also, there is a semantic thing. We have separated a logical division line. What we call PHY we can split wherever we want. We choose independent layer - means whatever PHY is there you can talk to it. That interface is just for ease of speaking - a conceptual interface. The real interface is the dependent layer. A nice concept. The MAC/PHY interface is logical only, it helps working in the subgroups.

The real PHYs separation is at the dependent layer. Also, now you have a management layer that you fit between the logical interface.

Francois: identify - call MAC/PHY interface logical as opposed to the electrical interface at the DTE/DCE interface.

Larry: but if you look at the ISO model or other standards, we are iterating toward this as the model. The ME is the media, so the antenna is at the bottom of the MDS and the exposed connection (if there is one) is above the PLCE interface.

Paul: assumes Larry is considering LEDs to be an antenna?

Larry: yes. Do we support where the antenna sits - on a MAU-like device (that's an exposed interface between MDS and PLCE)? That's not been addressed.

Paul: leave that problem to the people doing PHY layer. Our purpose here is to separate MAC and PHY issues.

Bob C.: why would you have a connector between MDS and PLCE? There can be multiple channels - we want the MAC to be able to control those channels (such as for redundancy).

Larry: down there the BNC handles bits - the PHY layer convergence entity takes symbols of the form the computer people understand and makes them into things radios understand.

Bob C.: in IR its easy to imagine individual transmitters that don't have a simple DTE - the whole point of this connector is to have a simpler connector where it is needed.

Larry: if you want to have a MAU, say so. If its an antenna that plugs on with a BNC, then have 5 BNCs if that's what you want.

KC Chen: we have a joint MAC/PHY meeting so we can proceed better in our separate meetings. We want to define the MAC/PHY interface - we shouldn't be discussing how to implement it. What can PHY deliver and what does MAC expect. We are not doing this.

Larry: we just tried to tie down architecture to help.

Francois: we have decided where the MAC/PHY interface is and this is all a PHY issue.

Larry: this is the architecture model. The MAC/PHY interface is logical concept. We should move on.

Francois: mechanical interface exposure is in the PHY in this model?

Larry: yes.

Michael Rothenberg: how will this architecture support multiple PHYs? What if I want to connect some kind of MAU to a single MAC? I have multi-connectors at the DTE/DCE connector. Now I have a single interface so I have a single PHY.

Larry: yes, there has to be PHY independent sublayer at the top of the PHY.

Bob C.: no you need another physical connection between MDS and PLCE. An issue should be recorded here.

Simon: given this model, if I have a given PHY how do I know from the MAC what I can control?

Larry: using the primitives that pass that kind of information across the MAC/PHY interface. A lot of that falls up there.

Paul: If we start talking about control lines at the MAC/PHY interface we can make progress. You have shown the separation is in the PHY.

Simon: suggest an issue - do we want a DTE/DCE interface in the PHY

Francois: there is an issue already - what is the MAC/PHY interface. This issue should be corrected to say what is a MAC/PHY logical interface. Then what is the requirement for a physical interface.

Simon: at the MAC meeting tomorrow, as a member, I would like to know what I can set in the PHY and how do I know that.

Jonathon: in the specification of what kind of PHY you have, the first thing is n-time redundancy. First there is redundancy. Then there is power level. Data integrity - how good is the data coming back. Receiver sensitivity. Health of PHY.
Simon: if I have a piece of data to send - on what basis do I make these choices?
John McKown: you would have a word of information about the status of assorted channels offered by the PHY to the MAC to support MAC decision making. The MAC could then make its choice.
Michael: the reason you want a MAC/PHY interface is to hide PHY specifics. Why pass them from the MAC to the PHY? This is why we have a PHY independent sublayer.
Larry: for example power control - in some proposals MAC decides by looking at receive power levels and transmit levels. The PHY has to tell MAC these things so MAC can make the decisions. What you say is wonderful for wired media, but not for wireless.
Michael: all this intelligence may be in the PHY.
Andy Luque: across the interface - this is a data interface. Maybe some status, but data up and down, the power sorts of things are management items.
Larry: we need the MAC to give us this information.
Paul: serial passing of information - how about two connections one being data and one being information.
Nathan S.: look at 92/61 - we have listed them.
Larry: they have been listed many times.
John: what are the lines? What interface are they part of - power control is management.
Larry: but power control on a packet by packet basis - this is not management it is data passing.
John: management is something that lasts for how long?
Larry: management is not something used to gain access to media.
Chan: some things like CRC are MAC things. Signal quality can be determined and should be in the PHY. Likes the model where a simple status, data and clock in each direction. There is a problem - management has to enter at a certain point but it must propagate. Controlled power at a station from another station - power control may have to go across the medium - multiplexing this information is difficult. The list of functions can hopefully be simplified.
Larry: we will not decode addresses in the PHY. Any function that requires knowledge of who or where, we will not do. You have to know how much power needed to talk to him.
Wim Diepstraten: functional approach, not implementation It is a logical thing. Let's first address what is functionally needed. Larry is also saying that we need that extra entity between MAC and PHY to allow management on a packet by packet basis. What kind of functions are needed on both sides?
Bob C.: reason PHY wants to present information - necessary for any type of media management strategy. Management as separated from MAC can intervene here. The redundancy manager in field bus intervenes between the MAC and PHY in real time in the same sense as these are being talked about.
Michael: good point about address decoding in MAC. But decisions about use of physical medium should not be taken by MAC. The MAC should pass to the PHY information necessary for a certain transmission, even on a per address basis.
KC: to help reach a conclusion, summarizes - (1) MAC should deliver to PHY data and CRC for sure. (2) MAC should send information to PHY about channel choice if there is channelization (whether for hopping or redundancy). (3) some PHY layer control - this is not clear but there may be something. (4) Some control of PHY - possibly MAC must pass address information to PHY. (5) MAC tells PHY procedures

to follow. That's five types of information MAC should send to PHY. We can discuss these five one by one - it might help.
Jonathon: Bob - redundancy in the model set out before, redundancy selection is made by MAC, whatever means used to pass that information down. Since MAC has all this information and PHY just carries it out. Control signals to PHY indicate what action PHY should take. Clean interface - MAC decides, PHY implements and returns to MAC information for decisions. PHY independent layer decides how to implement what it is told by MAC.
Andy: MAC should be independent of PHY. Right now we have a PHY box and a MAC box. If you try to but these functions into MAC they make the MAC PHY dependent. A generic MAC dealing only with data allows these PHY dependent items to be in MAC. What about the side management box?
Larry: PHY has to keep address related information and make redundancy decisions on this? We could just pass the information up to the MAC and have MAC use information it already has. PHY has informed MAC of options (diversity options) it uses this information.
Jonathon: premise that different PHY has different parameters. They can be grouped into a unified set and the MAC can recognize what PHY independent parameters to use. A set of functions is PHY independent.
Bob C.: agrees with Larry and Jonathon and Andy. Management can intervene on a frame by frame basis. Management can maintain the stuff Larry refers to. True MAC independency is through the management function.
Larry: then we should form a management group immediately.
Ian Crayford: not confident that the decision algorithm in the MAC will be identical for all PHYs. So a single MAC will have multiple decision algorithms dependent on PHY type.
Simon: we are mandated to one MAC, multiple PHYs. MAC must make decisions dependent on items - will they be the same for all PHYs?
Larry: turn power on, PHY says there are 2 redundancy options; MAC says ok; first message is received ok; MAC says PHY use option 1; third message comes back, and its no good; MAC can retry on 1 or 2 - he chooses 2 and it works; so now he knows that he got to that station using 2 and remembers that for next time he wants to talk to that guy.
Simon: that's one example. Power control - how do I manage that?
Larry: PHY says I have variable power levels, MAC can use them. If PHY comes back and says I don't have them, then MAC is going to not use them.
Simon: when I turn my station on, the PHY is going to tell me what my options are?
Larry: yes.
Simon: interesting.
Ian: is there a document describing this?
Larry: we are working on it.
Michael: what is the opposition to migrating functions to management?
Larry: MAC gets frame with bad CRC, he tells management to tell PHY that the redundancy option should be changed. Management tells PHY and PHY must remember that now. Assumes that all information must still be kept in PHY.
Michael: one normally detects channel quality before transmission. Make redundancy options known before transmission, can be done before the MAC knows there is a channel problem. You need source address information to make intelligent decisions. To make decisions in PHY you need MAC information and vice versa - so put all information in management and let it decide.
Larry: the MAC is the only one which ultimately knows whether the data gets through or not.

Paul: why in MAC rather than PHY - supposing the PHY detects power fall off, the only thing PHY can do is increase power. The MAC can decide whether to increase power or switch antenna. Has to reside in MAC because PHY can't communicate between entities.

Bob C.: station management can redirect media for every frame. It is off base to say it is too complex to do in management. Could we say either MAC or station management and move on with that?

Jonathon: optimization and redundancy - these are different from power control and antenna diversity - if the other side wants power lowering PHY cannot know. MAC must decode and tell PHY. That is optimum. All PHY knows is amount of energy - it doesn't know effectiveness of data interchange.

Michael: intelligent preamble can let PHY know those things.

Larry: receiver can choose antenna, but on transmission you have to rely on history - at the MAC level that history must be used.

Michael: why at MAC level and not a management level?

Larry: synchronization of events - how do you get the management coming in at the same time?

Bob C.: that's the implementers' problem.

Larry: on every packet the management entity is going to be rolling around collecting information in parallel to the MAC.

Michael: advantage - it divides problem into manageable chunks: PHY does things, MAC handles data; management optimizes.

Orest Storoshchuk: traditional MAC doesn't worry about management functions, but it doesn't appear here that the PHY can manage the medium. So there has to be a PHY management section. We are arguing should it be part of the MAC to interpret and control PHY management information, or we could have both layers communicate sideways to a management function. The main argument is how can you synchronize this information.

Larry: is sitting in the middle of the two choices right now.

Simon: a side management function may be the only way to satisfy the MAC and PHY people who both don't want to do this.

Bob C.: there are PHYs for which the side box will not exist, so there are practical reasons to do this. Limiting the scope of what we do first pass could be done using this approach, enabling us to get a standard out faster.

Wim: in this approach, the side bar is highly integrated and synchronized to the MAC. There will be a lot of duplicate of logic in the PHY in order to affect this.

KC: agrees with the idea. Problem - for instance, power control. One use is for MAC to use for cell isolation; another use, for power conservation. If purpose is for PHY, keep it PHY, if MAC keep it MAC. We cannot rely on the history of transmission, a reminder.

Larry: history is not for assurance but to increase probability.

Bob Rosenbaum: if we have a single MAC, it may be a huge MAC with PHY specific issues or a MAC with a management piece, either way we have just implemented separate MACs. Experience says that MAC management division didn't work. Maybe you need to narrow the focus and simplify things being addressed.

Orest: sideways management has synchronization problems. We could pass everything through management - make management a layer between MAC and PHY.

Jonathon: there is a unified set of parameters that can control all PHYs from one MAC. PHYs can act according to a fixed set of instructions. If you believe this, then the question is settled. The management side bar - you are trying to cut the MAC into little boxes. This management function is a part of the MAC. It has the information to make these decisions. Who cares where you put it - it is a MAC function and MAC has the information to make the decisions.

John: Bob, the media independent PHY sublayer - did you think of this as media management?

Bob C.: management comes in from the side. It is coordinated with data transfer on an octet by octet basis. There is a media independent layer between the MAC and PHY and a side management box beside the MAC and media independent layer (not down into the PHY).

Larry: what we call media management is media access control. Changing thresholds dynamically is that media management or media control.

Francois: to Jonathon: the little box in the MAC is connection management entity.

Larry: for time bounded services?

Bob C.: SMT does this, including media management.

Larry: this is a new thing.

Wim: we are identifying things, for instance power control, as media access functions. We still need a management interface directly to the PHY to accomplish this. This is a direct interface MAC/PHY, not through a management interface. PHY will have to duplicate logic to do that.

Simon: we are drawing conceptual boxes. The work has to be done somewhere. The box doing it could live right inside the MAC.

Larry: put it in there and take it as an action item for the MAC group. No separate management.

Simon: PHY must provide a media independent set of things that MAC can always set.

Larry: the PHY might give null answer when MAC sets one.

Michael: media independent functions differ now. Ethernet started with one medium and changed. Very different as each medium was added. It would be good for the sake of time to market and standard evolution to separate out to management. Don't care what you call them but they should be dealt with separately. We cannot go with abstract primitives referring to media quality.

unidentified: the PHY attachment interface should be defined first - the boxes are not as important as the interfaces.

Jonathon: it will be the PHY group's responsibility to come up with a set of completely independent parameters that the MAC can deal with. That set is the responsibility of the PHY group. The MAC just operates on this set of parameters. Algorithms might be recommended by the PHY guys, but how they use them will be up to MAC. But there is a set of PHY independent parameters passed to MAC, and MAC passes data and control.

John: doesn't like this - is there a standard MAC/PHY interface? Isn't that the set you speak of? MAC is ordinary MAC?

Jonathon: there is a MAC/PHY interface line. The MAC is not an ordinary MAC because of the things that need to be done.

Larry: PHY gives you a set of tools to use.

Simon: are you also going to say, given various circumstances you are going to recommend an algorithm for using them?

Michael: this abstract set of parameters will put us in a situation that will not permit us to evolve. We don't know what the reality is yet - what are the PHYs layers and how they behave. We don't know the strategies. We haven't experimented yet.

Jonathon: for the past 5 years and next 5 you could count the media on 5 fingers and know quite well the limits.

Michael: sun spots are changing, therefore so will our media.

<u>Jonathon:</u> doesn't feel the basics will change much - the fundamentals change slowly.
<u>Larry:</u> we are talking about trying to build an ordinary MAC. 802.3, 802.4, 802.5 and 802.6 exist already - there wouldn't be 802.11 if an ordinary one would work. 802.4L went 3 years trying to get an ordinary MAC. The MAC has to account for these things.
<u>John:</u> just because it doesn't, doesn't mean that's what you should change.
<u>Larry:</u> the MAC and PHY have to work together in a way they never have before. These aren't any old PHYs, these are state of the art PHYs.
<u>Simon:</u> a set of managed objects that are PHY independent - then it doesn't matter if the MAC or some management entity handles them, you have said the PHY won't do it. That set of objects - conceptually MAC or management, MAC group can handle that. That seems to be a concrete way to approach that - PHY group comes up with the objects and MAC is happy that they can handle them. I have a feeling the everyone likes that.
<u>Nathan S.:</u> we were trying to shoot at two targets and we need just one. There are too many good possible solutions. So choosing this one and working from it is a start.
<u>Simon:</u> maybe this won't work out as the right way to proceed - the set of managed objects eludes me, so I see a complex MAC - but you tell me there will be a simple set, so let's work on that.
<u>Jonathon:</u> currently we have MAC/PHY interface which will consist of a transfer of PHY independent control information. MAC people will work on what they can do with this. The PHY group has to come up with the set of objects. If we fail, then we have to try a different way.
<u>John:</u> if the PHY group makes the list, shouldn't it have a set of recommendations what to do with that? Doesn't that mean the PHY group is designing the MAC algorithms for how to use the entities?
<u>Jonathon:</u> MAC guys are used to doing MAC things, but now MAC guys have to do some PHY things. So PHY should recommend to MAC how to handle them. Since MAC has the information on how to twiddle the PHY knobs, they must do it. I don't care who takes the credit, but it has to be done.
<u>unidentified:</u> has been in a unique position looking at wireless applications, IS54 digital radio implemented. Don't hear anyone taking about control channel talk for this wireless LAN. If a station is required to lower or up its power it would just have the information, send it, pass up to the MAC and back down the PHY. An encapsulated control channel encapsulated in the data.
<u>Jonathon:</u> this is a different beast from IS54. It has related infrastructure we don't need. No LAN operator would conceive to have such an expensive cell site. This is a different situation. The issue of a control channel is a separate one not addressed yet. There are a lot of little issues that you have to see through - go back through the documentation.
<u>Simon:</u> in summary: if the MAC is going to have the option to manage management objects, the PHY must answer what is the set, if there is one. In the MAC do we have management entity in the MAC or separately, or do we believe that we can implement what we want with this set of objects.
<u>Francois:</u> conceptual model? Does the MAC or PHY group decide what is the service interface?
<u>Carolyn Heide:</u> it is still a joint issue.
<u>Larry:</u> and fairly non-controversial.

Meeting adjourned at 6 PM.

Thurs. AM, July 9, 1992

Meeting called to order at 8:30 AM chaired by Larry van der Jagt. Secretary Carolyn Heide.

Agenda:
- (1) Presentation of Input Papers
  - 92/76, Wim Diepstraten
  - off the cuff, Jonathon Cheah
- (2) PHY report
- (3) Open Discussion
- (4) Break, after which MAC and PHY will convene for brief individual meetings.

Some general announcements:

Vic - token ring standard still available at registration. Electronic copy of all submissions to Vic with headers and footers as specified in IEEE 802.11-92/00, or a hard copy with those, so they can be mailed out as soon as possible. Changes have been made in the ET NPRM reply paper, he will try to get copies out. Be sure to register (i.e. pay) if you haven't yet.

Dave Bagby - Sorry about not being here until now. Document IEEE 802.11-92/66 is double circulated to some people by accident. It explains some of the functional requirements definitions and things. The paper makes sense to those involved already, but takes two passes if you aren't - terms are used before they are defined. E-mail reflectors have been set up: double check your e-mail address in the attendance list.

Agenda Item (1) - Presentation of Input Papers

The Potential of Dynamic Power Control, IEEE 802.11-92/76, by Wim Diepstraten

Follow up of the Leiden meeting where simulation results were presented comparing distributed MACs.

Potential reuse efficiency by using power control is the subject of this paper.

When desired coverage size is larger than a critical area like - if you have small cells (femtocells) you may need at least 4 APs to cover a room. If the BSA is that small you also need a large number of channels to obtain sufficient BSA isolation. For cells which are large but bounded by walls, floor etc, then the channel sharing is not that bad. A double ring of cells is needed perhaps for isolation in first case between BSAs that want to use the same channel.

Figure 1, page 4 - the actual signal level would be the same curve shown but a thick line making an area around this curve. In a typical CSMA you have a typical defer threshold where you listen with a certain receiver threshold level. That is the defer threshold. Then in a homogenous environment (no obstructions like walls) this would be the defer average within that radius. The carrier sense logic can be more sensitive than the actual receiver. Given the specified SNR the inner circle represents an area where a receiver can hear from a transmitter at the origin. The second circle is where you would hear if there was no interference.

This is just an average - realistics controlling the fading margin are needed (later in the paper there is a picture of that).

This is of course only a two dimensional picture. This is a homogeneous environment which does not exist, attenuation boundaries like walls exist in the real world.

The double ring of isolation increases the number of frequencies you need dramatically - Wim thinks about 20 will wind up being needed.

Jonathon Cheah: the cell separation underlying this assumes that the installation will require a site survey - which is not desirable in LAN deployment.
Chandos Rypinski: disagrees - you might do a better job with the site survey - but you can do without it.
Jonathon: then forget about going to Egghead to buy your card!

This concept does not really have anything to with the type of protocol used - this is the SNR reality at the intended receiver to correctly receive your message - central or distributed protocol.

We have to reduce the co-channel interference level because that is the major interference factor. Reducing transmit power level can do that. A given power level allows you to reach the coverage area boundary. So we could do with less power. Must learn what is the proper level per individual station and maintain a table which identifies how much power is needed for each.

Bob Crowder: power level control is an alternative to having to go to the 20 channels?
Wim: no. Considering interference I drew a picture of what isolation is needed for bandwidth reuse of a channel. If you do not have the channels, then you have to share. However using power control you can reduce the number of channels needed.
Bob C.: so it's not an alternate to 20, but allows you to get by with 5 instead.
Wim: it cuts down the number of neighbors with whom you have to share.

A homogeneous environment gets a reuse efficiency of about 3. A typical environment - a few cells - and the reuse efficiency increases because you decrease your power to support only what you need - the cell size is dynamically dimensioned. This allows obtaining reuse factors much higher - 20 to 30 even. Summary - larger reuse efficiency as a function of BSA size.

Can operate in a mixed environment, so can be implemented as an option.

Larry: what would the PHY have to tell the MAC, and get told by the MAC, to make this work?
Wim: at a receive location what level to use to reach a destination can be learned. When a message is received from a certain source, measure the receive level, average it over a number of frames. When I know what kind of level, I know the attenuation path between us and could use that knowledge to reduce the level to a level sufficient to achieve a sufficient receive level for a proper SNR at his receiver. A transmit power level indicator would be required in each frame. Also the silence level (which is what Wim calls the interference level) would be needed in each frame in case there is an interferer in the area of the station, making the level at that station deceptively high - an interference level which is none co-channel. Two bytes in frame structure for these two things. In the PHY a variable gain output stage and a receive level measurement function and if CSMA a variable defer threshold. In the MAC you need dynamic control of transmit level and defer threshold in a table for each possible destination. Also a receive silence level. See figure 3. Figure 4 shows what is needed from a more implementation point of view.
Bob C.: does the receiver have to know the silence threshold before it knows whose frame it is?
Wim: the silence level is monitored on a regular basis not per packet.
Bob C.: depends on various nodes doesn't it? Depends on who is transmitting to me?

Wim: no, you can monitor in an interference space when everyone is silent.

Certain implementations may not have an receive level measure or variable power. Could still work provided proper transmit level is indicated in the frame - if the guy you are sending to is not using power control you can still use it to send to him. Any implementing station benefits.

unidentified: what if you're in a null point - couldn't you be deceived into thinking you need to blast power?
Wim: a turbo mode is not provided - you reduce your normal level, you don't increase.

Simulations: using simulator described in 92/26, with the power control options. Network with individual stations 1 thru 7, number 0 is the server. There are peer to peer tests in which every station transmits to a random destination, and client-server environments like Novell perform 3 (all under high load) where all station transmissions go to the server.

Note that although WAVELAN is one of the protocols simulated it does not have this power control implemented.

Jonathon: CSMA-WAVELAN performs better that the other two. Why?
Wim: there is little overhead in that as compared to the others.
Payne Freret: why is the performance less in client-server simulations?
Wim: this is explained in paper. The perform 3 test does not include MAC overhead, the other does.
Nathan Silberman: what about delays reading the extra power information?
Wim: no, doesn't required medium time.
Mike Bergman: how is defer threshold learned?
Wim: start with maximum power and learn eventually.
Mike: but it varies from destination to destination.
Wim: yes, that's why you maintain a table.
Bob C.: how do you learn - the far node sends you the power at which it transmitted. Do you assume that if there was no CRC error its ok?
Wim: no. What you learn is what the attenuation of the path is. You will need to use some type of translation of that to obtain the power level to which you could reduce for the destination, to which you can assure a certain SNR at that destination, which also has to do with the carrier sense you use. You measure your receive level per link. From the transmit power you know the attenuation so you can learn the average - add a fading margin and translate to a transmit power.
unidentified: benefit is in a dense environment. How do you prevent from system from chasing power up?
Wim: you don't go up. You start at your typical power level and hope to reduce.
unidentified: you are going to have high interference levels in this environment. It would keep going up.
Wim: you reduce such that you assure a sufficient SNR at your destination. You don't lower until just above its receive threshold.

Wim has just explained why power information might have to go across the MAC/PHY interface on a frame by frame basis. Now Jonathon has a proposal for how information could go over the MAC/PHY interface.

MAC/PHY Interface Information Transfer Concept, Jonathon Cheah

Jonathon presented the following 4 hand-drawn overheads (sec note: which were later copied and passed out to attendees, handwritten and without a document number. So I will reproduce them here for non-attendees):

Slide #1:

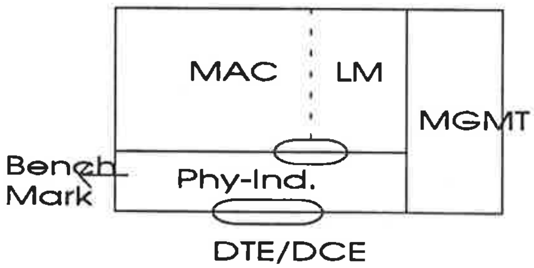

General model concept

a, b, c, ... from Higher layers

MAC

, , , ... from Phy.

f(x), f(y), ...

Algorithmic concept:

f(x) a control function from MAC to Phy

where x = f(a,b,c, ..., , , , ...)

for instance:

x = a c

Slide #2

A. Primitive: from MAC to Phy
1. MAC_Driven_Phy_Passive (Channel, Options, CAT)
2. MAC_Driven_Phy_Auto (Optimize, DOG)

B. Primitives: from Phy to MAC
1. Signal Quality
2. M_D_P_P_REP (CH/option)
3. Phy_generated (health + status)

(C) Parameter Independent set
1. Phy_Kill <panic button>

Slide #3:

CSMA

MAC_D_P_P =
- a. freq 1
- b. freq 2
- c. freq 3
- d. .

f. power level 1
.
.
.

IR

MAC_D_P_P =
- a. Diffused

b. Focused
c. Array Sense
d. Single Sense
e. Section Sense

Slide #4:

SALOHADAMA

let MAC_Driven_Phy Passive =

a. code 1	f. power level 1
b. code 2	g. power level 2
c. code 3	h. receiver sensitivity
.	i. receiver sensitivity 2
.	j. receiver threshold 1
.	etc.

let MAC_Driven_Phy_auto =

aa. Antenna Diversity
bb. Equalization
cc. ... <any other Phy optimization>

<data><Signal Q><M_D_P_P_Rep><Phy_gen>

Presentation and Discussion:

The divergent PHY concepts from the group may be leading us down the wrong path. We should isolate the MAC parts and the rest are a PHY specific download parameter set. This presentation goes quickly through the mathematical algorithmic concept.

Take MAC as a black box - it is actually a state machine that performs functions according to external stimuli. Given parameters to be determined by PHY group or MAC group - then there is a set of parameters that come from the PHY. MAC state function takes "" and does a function of some kind - x, y or z.

Given that we have a set of variables each of which can be defined separately, the idea of an independent MAC becomes possible.

For example: on slide #2 A.1 says optimize transmission which causes, A.2 antenna diversity - which could be autonomous - PHY could find the best antenna itself. On that slide (C) is the only absolutely mandatory thing Jonathon could think of.

Slide #3 shows how the PHY designer defines configuration. Common types of things that MAC reads off - say you have a primitive that has 100 variables, I can set: send a = set code 1, b = set code 2, etc...

On slide #4 AA enables PHY to take action autonomously.

MAC says execute a, b and c. MAC knows from initialization what functions to do if a, b, etc.

If the packet quality is consistently low and the upper layer is upset about it, and PHY says I optimized and cannot do anything more, then MAC says lets negotiate a change of code. If that is function x, then MAC starts to execute that function.

If MAC knows from the initialization process what these parameters mean, then there can be pre-programmed responses.

Physical layer designer defines these tables as seen on slide #4. If you build any PHY you can define the algorithm - the action taken by MAC - which may be completely different per each PHY. This allows all PHY manufacturers to develop the best they can and still maintain a single MAC.

Problem: there are sets of parameters that wouldn't make sense. Say IR in a room without too much mobility, the BER is 10^{-7}, and the best case is 10^{-12}. But for CSMA it could be that 10^{-5} is great and 10^{-3} is average. As far as MAC is concerned since it has to be able to give a number and do a reaction, there must be some bench marking system. Say if the BER is 10^{-5} in the IR you might want to take some action, but in radio 10^{-5} is great so I do nothing. So how can you be able to decide which level - there certain bench marks, say mapping a level to a number. In IR you might say if you reach 7, you might go for 9 - you optimize to a certain level which you think is the best you can get for your media. This PHY independence is a quality control interpreter. The PHY independent layer should be able to set the quality you want - it tells MAC the best you should want is x. These are sets of parameters that you want to normalize for all PHYs in the PHY independent layer.

Say for instance - although this doesn't need to be so - the PHY can introduce an independent function which is variabled. Therefore you can even put on the MAC control function x "if you see a 2 you are in good shape, do nothing". This is very much a PHY dependent algorithm. When you do a PHY you don't need to fill up all the variables either - you put in a no-op and when MAC reads the initialization file it knows to ignore that parameter.

If this is acceptable it is only up to the committee to decide how big the primitive is - we want to get a length of primitive that is long enough to satisfy all, and add 2 or 3 extra for luck.

Summary: Jonathon maintains that by using this definable primitive between the MAC and PHY we get a very standalone MAC and satisfy all possible variations of PHY imaginable.

Orest Storoshchuk: Wim said that MAC has its own algorithm to handle power - how does this fit?
Jonathon: say f(x) is the power control algorithm and is the noise threshold. "a" is MAC preset power level to maintain. So when noise threshold goes up, PHY has defined f(x) to say that given the power level and noise threshold I should give command "g" - set power level l2 - and a command "h" to give receiver sensitivity 1. That's it. The PHY knows "h" and "f" mean do such and such a thing. Give that same command to another PHY and it would be completely confused. But the 2 PHYs have completely different initialization files, so this wouldn't happen. Initialize with the wrong parameter file and you're screwed.
unidentified: MAC would still have to implement things with knowledge.
Jonathon: the isolating S/W or F/W is supplied by the individual PHY manufacturer. MAC is an empty state machine for individual PHYs.
Joseph Mazor: there are already TI chips that have downloadable S/W. There are companies that download better S/W and get better performance, but this is implementation. This proposal gives an abstract framework for passing information, whether the state machines are custom, given by PHY producers, or gradually defined by this committee. I applaud this proposal because it may give us a very simple but powerful framework to quantify the methodology. Nothing WLAN related - just a framework, information independent.
Jonathon: will submit this as a proposal.
Larry van der Jagt: we will figure out how formalize this.
John Deane: any clues as to size of configuration files expected?
Jonathon: depends on PHYs - for a CSMA, maybe 10 octets.

Bob Crowder: applauds this concept, but one more level is needed. Things that the MAC can take similar action on have to have the same name so the MAC can have a common set of logic.
Jonathon: personally, is sitting on the fence on that. It puts the onus onto the PHY F/W - there are more degrees of freedom.
Joseph: let's adopt this thing first and try that inside of this framework.
Steve Chen: Manufacturer interoperability is a requirement.
Jonathon: this doesn't preclude that. Isolate the kernel of the MAC, the downloadable part is the option.

Agenda Item (2) - Summary of PHY group primitives, by Larry van der Jagt

Larry presents the following overhead:

All PHYs shall support a single channels. Support of additional channel is optional. If more than one channel is implemented a method of informing the MAC about the number of channels, the channel in use and to allow the MAC to change channels will be implemented.

All PHYs shall support a single level of transmit power. Support of additional transmit power levels is optional. If more than one transmit power level is implemented a method of informing the MAC about the number of levels available and the value of these levels and to allow the MAC to modify these levels will be implemented.

All PHYs shall report the status of receive signal strength relative to one threshold. Additional thresholds are optional. If more than 1 level is supported a method of informing the MAC regarding the number of levels, and the values of these levels will be implemented. The indication of receive power will be delivered on a frame by frame basis.

CHANNEL - an instance of medium use for the purpose of passing PDU's that can coexist with other instances of medium use.

Example:

single channel	n-channel
1 narrowband channel	FDM channels
DSS with 1 code	DSSS with CDMA

PHY group tried to come up with the set of things that the PHY will supply to the MAC.

Implication - if there is a MAC that needs a channel for information and a channel for data, then it will not work with all PHYs - one channel PHYs must be supported.

Power level control is optional. There will be power level and channel availability reported to the MAC, but MAC must accept not changeable as the response.

Signal quality indication - there are many things that are involved in signal quality. We will do everything we can do deliver a good signal. MAC will be able to tell PHY to try things on a per frame basis - signal quality will be reported on a frame by frame basis.

Bob Crowder: receiver level, or additional information?
Larry: This is only an example of where we're going. We haven't put it together yet. The most significant thing is shall support one channel. This is cast in concrete.

Agenda Item (3) - Open Discussion

Chandos Rypinski: caution for group on the subject of special PHY requirement stuff - it remains to be seen whether there is no alternative to parameter adjustment. Don't close your mind on other tradeoffs to simplistic systems against all of these parameter tradeoffs between MAC and PHY.

Bob Crowder: you may be right. But the way to come to that is to say that for that simple PHY there is only one or even zero parameters. But we keep having presentations and we can't understand what they mean, but when tabulated by parameters they will become understandable.

Dave Bagby: where intelligence lies with respect to the MAC/PHY boundary - Passing this intelligence across this boundary implies intelligence on both sides of the boundary. This seems to be opposed to the principle of keep PHY simple and let intelligence stay in the MAC. Was this discussed?

Larry: if the PHY ever has to know an address we have gone too far. Or if it ever has to decode an FCS. PHY will never decode data - any function that requires that is not allowed.

Mike Bergman: in response to a primitive to twiddle knobs - could the PHY respond to a command for something it doesn't have? This avoids downloading. PHY just says no.

Larry: I think of it as uploading from PHY to MAC on initialization.

Dave B.: discussion of how to determine what the other side does is premature. Concept of knobs and trying something different on each frame implies protocol. Twiddle a knob on data frame and you have no way of knowing whether it will make things worse. A ping - a dummy frame that just tries changing things instead of wrecking data - did you discuss that?

Larry: recognizes that it may hurt instead of help.

Dave B.: mobility implies that dynamics may change before the right knob setting can be found. We may wind up not taking advantage of what we've got.

Jonathon Cheah: the model is generic. PHY autonomous primitives allow different PHYs to handle things different ways. Management knows what actions are best for PHY so can set autonomous things appropriately.

Jim Schuessler: we're trying to talk about specifics. The general issue is self configuration verses apriory knowledge. Both should be accommodated. There is an issue of how this configuration information is obtained - self configuration or user configuration.

Larry conducts straw poll - how many people think Jonathon's proposal sounds good so far - at least 2/3.

Joint MAC/PHY Meeting adjourned at 10:30 AM.