Tentative MAC Minutes
Tuesday PM, November 10, 1992

The meeting was called to order by chairman Dave Bagby at 1:40 PM. Carolyn Heide secretary.

Subjects for MAC group consideration this week:

- time-bounded services
- MAC/PHY interface (already covered in joint MAC/PHY meeting this AM);
- distribution system.

IEEE P802.11-92/129, Time Based Services QOS Requirements on a Wireless LAN
by Bob Crowder

Most of this is in 92\101, this is an update specifically on issue 15.1. In the functional requirements process we identified 4 types of time bounded services, which are well described in Tim Kwok's paper.

All of these processes deal with the reconstruction of natural processes, e.g. speech. All of these add dimensions and bandwidth requirements to the reproduction of the signal. What are natural processes - movement of objects producing sinusoidal waves. The real picture is lots of harmonics superimposed on the media. Reproducing all the significant frequencies is the important thing. The Shannon theorem says that you must sample twice within the highest harmonic. In industrial application usually between 5 and 10 times more than that within the harmonic is used.

A fundamental requirement is that the variable phase shift doesn't vary by so much that you can't reconstruct. (jitter). Feedback loop places additional constraints - you can't compensate for jitter because the control system has already compensated.

Quality of service (QOS) - Bob believes that there are several QOS parameters (in the OSI sense of the term) important for any time based service. These are similar for speech and industrial.

Data rate requirement is fundamental. The difference between broadcast and freeze frame video illustrates the wide range of requirements. Industrial rates are typically lower than speech rates but very variable.

John C: how many industrial applications are digital already?
Bob C: probably over half.
John C: the bit driven requirement where a switch closes and you have to report this within a certain time - no sampling, but a time bounded event is generated. How does that relate to this?
Bob C: you would have to allow the event to come within the natural frequency which puts a higher load on the system than the sampling.

The real point is that all of these numbers have the limits. Mechanisms that guarantee 64 kbit/s for speech will be able to guarantee the industrial control margins very comfortably.

Wim: the total spectrum is from .1 to 100 samples per second.
Bob C: there are samples per hour, or per ms, this is not the total spectrum needed. No one could safely state the outer band in industrial control - just try satisfy nearly all of the market.
Wim: with 100 samples pre sec you could say this cover what % of the market?
Bob C: a very respectable % of the market.
John C: this same category covers many other things too, like patient monitoring in hospitals. The next issue is not how much dedicated bandwidth, but how much time you can allow slip.
Bob C: it is typical in industrial control that you don't transmit just one value, you transmit something like 8. So at the sample rate you could have larger packets
Wim: is looking for some understanding of what the spectrum of requirements are and how they relate.

Second QOS requirement - maintenance of the cyclical nature. Must maintain a pattern that corresponds to the frequencies in the natural process.

There has been a huge amount of study of voice sample and reconstruction. There are millions of industrial processes and the time has not been spent to study them. Much less precise definitions are available than as for speech or even video.

Third requirement - jitter between transmit opportunities. Speech, industrial control, patient monitoring all require this to be limited.

Fourth requirement - discarding of previous data.

Fifth requirement efficiency - there is a phenomenon in speech where the line goes dead (pauses by speakers) ,recapturing this bandwidth would be advantageous.

Sixth requirement - call duration.

John C: in any of these cases, are any of these systems appropriate for a contention based connectionless service vs. a dedicates service?
Bob C: trying to present QOS parameters that 802.11 needs to present to make it appropriate.
John C: this addresses a lot of circuits. Call duration for instance - this is not contention oriented. Most industrial controls are circuit based, not connectionless services. Is there a class of time-bounded services which are appropriate to the connectionless service?
Bob C: all time-bounded services are connection oriented because of the call duration nature.
Dave B: there are other issues about connection versus connectionless service. Here's some way to get some handles on some of those services, this helps categorize them.
John C: a lot of these are simplex. Speech is full duplex, video is simplex. What is the impact on mobility - how many of these services move or are fixed?
Bob C: is trying to say if we provide these services, this is what we should do. Since we provide mobile services, we will.
Dave B: if we find we have said we will support contradictory things we will handle it at that time.

Speech, video and industrial control share the same parameters but with different values.

Wim: doesn't agree. When values are off by such a large magnitude, it becomes a different case.
Bob C: but these are things are the things that must not be ignored, regardless of value. They must be considered. Bob feels that mechanisms that support speech will support industrial automation.
unidentified: what about the role of lost data - speech can survive lost syllables, industrial automation can't.
Bob C: broadcast data can be used in an industrial control system.
John C: the bandwidth requirement for speech is similar, but the emphasis for industrial control may be on bit error rate or jitter. Also on the low end it's silly to discuss time bounded services for 3 bits per month.
Bob C: if this committee does its job right, if we support speech we will be OK for industrial automation.
unidentified: what about the lost data question?
Bob C: can be protected above the MAC level.
Dave B: MAC doesn't generally provide guaranteed delivery. This may be a major problem for the industrial control application. If someone above must have time to receiver, then that must be factored into the delivery time -this is not easy.
John C: session level control is requirement for this end to end ack.
Bob C: this is resolved at the application level in field bus, there is a ticking clock that says I should get the signal every 20 ms, and if it doesn't appear a couple of times in a row human intervention is required. Moving it out of the MAC is very practical.
Wim: dropping late data in the MAC? How is this done in industrial applications?
Bob C: enforced out of the system - the MAC wouldn't transmit the frame in the first place if it was too late. this is a transmitter problem.

Wim: if there are store and forwards nodes, where does it drop?
Bob C: at every forwarder.
Dave B: please let François know about relation to issues so that we could mark the document number on the issues log.

ATM Within 802.11 MAC, A Good Idea? (issue 20.3), by Jim Schuessler

Contentions:

1) TB services must have a guaranteed bandwidth mechanism.
2) ATM is one such wired (fiber) LAN structure meeting this.
3) We must be able to work with ATM backbones.

Conclusions:

- Use of specific ATM-MAC is an inefficient and unnecessarily constraining structure for WLAN.
- Present frame based (fixed length) MAC proposal meets TB service needs.

Proposal: Use ATM as a sub-MAC frame instance which has a variable bit/sec rate.

Conclusion: not a good idea.

Advantages:

- BW on demand - more efficient for MPEG for instance which has a variable bit/sec rate.
- maps easily to future WLANs. (B-ISDN based)
- It's modern.

Disadvantages:

- ATM design based on two assumptions which are not true for WLAN:

1) BW is plentiful (i.e. we can sacrifice some efficiency for self routing characteristics).
2) Channel is reliable.

There are examples of "modern" MAC's based on a fixed length, 8 kHz locked, frame structure:

1) P-1394
2) ISDN
3) FDDI-2
4) ISO-Ethernet (proposal for new 802.9 PAR)

The main point is that the ATM frame is too inefficient for use on a wireless MAC because we need to put our header on top of theirs. Further than that, it adds complexity to constrain time-bounded services with a variable length frame - it doesn't work well within the superframe structure.

Discussion
John C: the protocol or frame structure is not suitable?
Jim: the concept of bandwidth on demand is not suitable. All of the MAC proposals which say they support time-bounded services meet the time-bounded services needs.
John C: the time value of data is not part of the MAC proposals I have heard. There is just a retry life.
Dave B: doesn't believe we have addressed the issue of data life at all. obviously we do need to consider that.
John C: no MAC proposals have been connection oriented.
Jim: we have reservation based protocol, and we have slots reserved for time-bounded services in all proposed MACs - these could be used to set up a connection.

(Group) Random Address Protocol, IEEE 802.11-92/131, by KC Chen

Considerations for the MAC (remember the 21 criteria) - this MAC meets the following requirements:

-throughput and delay;
- multiple PHY support;
- seamless service: operate in multiple cells rather than single cell;
- time bounded services: unclear on the exact definition, takes meaning to be can't afford a very long delay;
- fairness;
- power consumption: for mobile stations;
- simple to implement.

KC proposes: the basic idea is considering the environment for mobile nodes in a multi-cell environment efficiently and reliably. There is the centralized and the de-centralized approach. Polling is centralized, but this proposal borrows a CDMA-like concept to make it un-centralized. Refers to it as a WLAN with infrastructure. We have a backbone network with base stations connected, and possibly repeaters connected to the base stations. For IR we must have the infrastructure, and for large scale radio networks.

Wim: the slide shows different cells - are those different channels or all one?
KC: is assuming one, for IR we prefer one channel although more is possible.

Approach: we have been assuming that in a centralized system the base station needs to know all addresses of stations under his coverage. This is not easy, if not impossible. It is costly in terms of bandwidth and efficiency. Traffic can be divided into two cases - uplink (to the base station) and downlink (from the base station). For downlink we have to use broadcast (in reality everything sent on the downlink is always broadcast if you think about it!). If there are multiple base stations, they will broadcast in turn, not simultaneously avoiding co-channel interference. On uplink, it may vary - ALOHA with low efficiency; increased efficiency with slotted ALOHA; carrier sense maybe better yet.

RAP has mobile units deciding their own "addresses". The only thing the base station needs to know is what the active node(s) is - i.e. those that have data to send. Why monitor those that have nothing to say. So the steps in the superframe are:

1. Base Station broadcasts [READY]

2a. ACTIVE mobile node(s) generate random numbers.

There will not be many of these nodes with simultaneous need to transmit so the dimension of this number need not be large.

Wim: all random number transmissions overlap in time?
KC: CDMA can be used to simultaneously generate these packets; it doesn't matter.
John C: any active node can select one of 'n' transmitters, and in the base station there will be one to n receivers.

2b. All active nodes simultaneously transmit their own random addresses

3. Base Station detects multiple addresses

4. Base Station polls according to the (chosen) random number.

John C: random number selects the frequency and the address.
KC: a frequency component of channel, or whatever, and the address.

5. acknowledgement.

Two units could try to use the same random numbers could be generated, so acknowledgement is still required.

unidentified: hidden transmitter problem - two units generate the same number, transmit, base station sees only one, sends data, both receive that data.
KC: acknowledgement is tagged to the data, not random number.
John C: 3 wire polling system is analogous.

Simulation and analysis data (see the document) is presented.

Tim: node in overlap of cells - do they register to a base station?
KC: you can respond to any base station.
Jonathon: What if a base station from a different network overlaps - these base station would not be synchronized.
KC: broadcast of ready packet identifies the base station - the receiver can identify if this is a base station on its network.
Jonathon: but you rely on time overlaps, so collision can be caused though.
KC: yes, but this is everyone's problem, not specific to RAP.
Tim: frequency of ready's - this becomes a function of the number of base station in the network as they pass control around in a token fashion
KC: if they know enough about topology they can transmit at the same time if they don't overlap. This is the same for all protocols not just this one.
François: assumptions are being made about the distribution system?
KC: yes. Assume base stations have knowledge of each other. You can make it as simple as you want - time token alone would do it, like token ring. The overlapping cell problem is generic to all systems.
Wim: multiple receivers - what is the dynamic range over which this concept can function?
KC: haven't gone through all kinds of simulations yet. But this is similar to the CDMA problem, so perhaps we can only deal with maybe 10 or 20 dB dynamic range. A capture-like phenomenon will occur.
Tim: downlink? Could it wipe out the adjacent cell.
KC: it always broadcast, always done in the broadcast period of the superframe.

On the simulation graphs. overhead is the ratio of overhead to data. So 1.0 means for every second of data there is a second of overhead.

Jonathon: why is that knee in the curve?
KC: probably a resolution problem.
Jonathon: looks like CSMA, so as you increase an asymptotic degradation will be seen.
KC: yes. But the CSMA curve is much sharper. We rely on reliable polling, not the carrier sensing.

How to achieve peak performance? One possible problem is what about one cycle with lots of traffic in it. This is addressed by the 'G' that turns the RAP into the GRAP. Traffic can be grouped so each broadcast ready is aimed at a group. All groups a period for time bounded services. There can be a group for new unit registration too.

Downlink traffic sent to units via their wired network addresses. These go sequentially in the broadcast period. They are acknowledged, but it is nit yet specified whether the ack goes in the broadcast section or in the uplink traffic.

John C: can multi-channel PHYs be used with this technique?
KC: sure, best to use extra channel(s) to help with the cell overlap problem.
unidentified: how does a base station find a station for downlink?
KC: there can be a cell hand-off.
John C: who knows in what cell a station can be found?
KC: the paper describes the algorithm for this. It is done like a telephone call. The hand-off and forwarding algorithms can be specified. The station plinks a packet to tell the new base station that he has moved. Store and forward is performed on the backbone.

Jonathon: as long as the mobile unit sends "here I am" every now and then this is no problem, and it is the same for any protocol. Time expiration is required on that, drop a unit if not heard in a certain amount of time.
KC: does discuss this in his paper because this protocol may make it a bit easier.
Tim: frame structure - group's are within one cell?
KC: yes. Say there are 100 nodes in 1 cell. If last time I used random number 3 successfully then I am in group 3. The reason for this is to further avoid congestion - it gets distributed over the groups.

This protocol uses bandwidth very efficiently and only requires one channel, although the number of channels used irrelevant.

KS Natarayan: have you done the 21 point comparison?
KC: not yet, but will do this.
Dave B: base station co-ordination makes cell overlap a problem. Unconnected base stations can cause collisions to a great extent. There is no mechanism to deal with collision because you assume the base stations are co-ordinating.
KC: counting on a certain degree of spatial isolation. Overlap of unconnected base stations will cause degradation, but only for certain cells.
Tim: dependent on transmission speed - 1 Mbit/s IR, for example. How can you afford all this polling time and the token passing time on the backbone?
KC: you can optimize in hardware.
Tim: the number of active stations effects how long it takes the token to get to the next base station.
KC: the superframe does not have to have a fixed length - if there is no traffic the superframe can be very short. You can spend varying amounts of time in each cell.
Carolyn: but it could have a maximum.
KC: parameter setting to stabilize and optimize is important, and flexible.

Meeting adjourned: 5:05 PM.

Wednesday AM, November 11 1992

Meeting called to order at 11 AM, by chairman Dave Bagby. Carolyn Heide secretary.

Distribution System Services Functionality, IEEE P802.11-92/128, by Dave Bagby

At the last meeting we talked about the distribution system (DS) and its services as part of trying to decide whether we need to define the internal workings of the DS or just the set of services it provides. Feeling seemed to be that the set of services was the right approach. This submission talks about the general categories of service - what should it do for you. Obviously it must get data to its destination, but is the only, or even the most important function.

This is not intending to imply that DS services (DSS) have to be implemented in an access point (AP). APs make an easy illustration of function.

Jim: distribution system medium interface (DSMI) - is that at the boundary of DSS to AP?
John C: is there a term in the requirements document now (DSMI)?
François: no. DSM (distribution system medium) is though.

Slide 5 - somehow a mobile station (STA) is associated with an AP. Reassociation - some way to change that, handoff or roaming, dynamic. Disassociation - a way of getting rid of these associations. Distribution - a way to distribute thing, or those things that deliver MSDUs within the distribution system. Making use of the association information is the key to this function. We have never been concerned with how this gets anywhere, simply that it does have to get there.

Slide 6 - authentication and registration have been used previously. Dave thinks there is little or no difference between these. They are related to security. It is the action of - "I tell you who I am" and "you decide whether or not to believe me". How strong should this be? That is a separate subject. A mutually acceptable level of authentication must be established. This could be as simple as just, "OK I believe you", or others may want tight authentication. Privacy - encryption perhaps. The thought is that it prevents unintended readers from getting information from your messages. Again you need a mutually acceptable level of privacy. 'Mutually' means STA to STA, more than STA to AP. Integration (Dave doesn't like this choice of word) takes traffic that gets into the DS and enables it to get into existing networks. Why is this any service at all? It could be that you just use your existing network for this, but maybe you need something special This just allows a service for that in case. Network management - Dave doesn't know what it is but knows that it needs to exist.

François: authentication, privacy and integration - this is called the DSS. Aren't these function 802.11 oriented? the DSS can be any existing network, so do these functions exist only in the AP. Because authentication and privacy may exist if the DSS is another LAN or WLAN, where do these functions reside?
Dave B: hears 2 questions here. Where, and how interact with exist things. Mutually acceptable levels if in an existing networks it is not mutually acceptable you have something to sort out. mutual is fuzzy whether AP to STA, or STA to STA.
John C: AP refers to AP on a BSS connecting you to the DSS, not the portal, right?
Dave B: is referring to authentication and privacy there, not through the portal.
John C: farther afield, authentication with the STA at the other end is needed too.
Dave B: yes there are two levels.
Phil: agrees with François. These are 802.11 functions not DSS functions.
Dave B: the DSS is part of 802.11 as per the par. Privacy in some sense affects the traffic in the BSS (air) but to create a mutually acceptable level some interface is needed that isolates what goes on in the DSS. Where does the DSS live - many places. If a DSS had a single computer then there is the DSS. Or if every AP is a CPU with a WLAN interface card and one to the LAN, then the DSS lives in every AP. No packaging of the DSS in AP should be implied here, just a boundary from wireless media to wired media.
Jonathon: we are talking about MAC. Authentication is above the MAC. The sense is that in the BSA you want membership kept to those legally allowed to use the service. This implies a level of complication which belongs higher up the stack. An encryption key could be given to anyone belonging to the system, there is then no need for the AP to keep registration information.
Dave B: what is your question?
Jonathon: what is the benefit of putting authentication into the AP? There are simpler methods to do so.
Dave B: is not claiming that it has to be there. But thinks Jonathon really means what is the benefit of having an authentication service. This service does not happen in the MAC. But we need to say that this facility must be available to be invoked. These services must exist somewhere outside of the BSS and we must be able to find these services.
Jonathon: authentication is not the right word then. Isn't this an encryption key?
Dave B: during the presentation please ask for questions for clarification - disagreements save for afterwards.
Bob C: integration - is this what we might call a localization of service, or the service that tells the portal which portal 'owns' a roaming node?
Dave B: yes, no ,maybe.
John C: one reason to control access to a BSS is that limited frequency resource means you need to know if you want to service any new node in a BSS or look at him as an interference source. If you have an infrastructure managing a domain - identifying who is in what service area, that process has to be part of the system
Tom Phinney: authentication and privacy are 802.10 functions and maybe we need a liaison meeting. This is not a subject for this group alone - we infringe charter and we are not the experts in this.
Dave B: is taking the middle ground -some interaction is needed, but we need to discuss it more to determine that.

Slide 7 - what is the life of a STA? To get associated with and AP you invoke 3 of the DSS services - association, authentication, and privacy. Authentication is not essentially part of 802.11 but the STA may need a way to invoke something.

Wim: authentication and privacy - related to 802.10. Is it local or end to end?
Dave B: right now talking about STA to AP. To become a member of the BSS, not whether the STA can get out of the BSS.
Wim: mutually now means STA to AP?
Dave B: yes. The reason I didn't say where to where is because an AP is an STA, so this could grow to cover the ad-hoc STA to STA.

These 3 things link the STA to the AP. An assumption has been made in the past -address is location, if you send something to address 5 it goes to him. This is not true for 802.11 . The association of an STA with an AP facilitates this.

Slide 8. - Association is the normal working state. Reassociation is how you know when an STA moves. These concepts have tended to get involved in MAC submissions, and we need to step back and evaluate these as functions and how does each MAC meets them.

Tom T: is there distinction in association state between talking to an 802.11 STA and a non-802.11?
Dave B: if I am a mobile node I am associated with an AP, that does not mean the AP is the end point of the MSDU. That is separate from the association.
Tom K: right, so why make the distinction between 'and other station' and 'and existing networks'.
Dave B: you get things into a DSS, Dave is trying to say there were two different ways to get it out - an AP or a portal.
Tom K: once on the DSS you don't care about the destination. But what about using an AP as a repeater?
Dave B: a BSS has a physical range but you have direct communication. A DSS is to extend coverage. Is a repeater a DSS? Could be yes, but it is not an alternative to a DSS.
Wim: how do the other services link to this - where is authentication and privacy used here?
Dave B: to create the association you used authentication and privacy. This is where Reassociation gets complex - you have established a working process with one AP, how do you recreate this with the next.
Bob C: authentication is a quality of service (QOS) parameter. When you establish an association you want to choose your level of association. This could be a QOS parameter that you negotiate when establishing an association.
Dave B: somehow you decide on this mutually acceptable level, and you don't talk if you can't agree on the level. A hook must be left to find this information.
Bob C: if a QOS is ever going to be needed, even once, then it must be tabulated, with null values if necessary.

Slide 9 - Dave believes the definition of ad-hoc and peer to peer are different. The former happens if 2 (or more) people walk into a room and start an information exchange. The latter implies something about upper layer functions, resource allocation in a networking sense.

The scenario becomes more complex with more than 2 in a n ad-hoc scenario, this is not reflected in these statements. Do you have 2 association, or one complex 1? Dave hasn't sorted this out, it needs more work

François: have you thought about whether Reassociation is null for ad-hoc networks? And what about roaming from one ad hoc to another?
Wim: an ad-hoc is a single BSS ESS. There is no roaming between ESSs.
Jonathon: ad-hoc network is a BSS without a backbone connectivity?
Dave B: yes.
Wim: each STA looks like an AP to the other STA? what does that mean?
Dave B: basically this means that it does the same things it has always done to associated, without being aware that this is an not an AP

Allan: ad-hoc is not the same as infrastructure scenario. In the ad-hoc each must establish a connection with the other but the other does not. Who has the led role here? To whom does the 3rd make its association?
Dave B: personally, Dave thinks it's a matter of who speaks first.
unidentified: can a portal be a member of an ad-hoc network?
Dave B: no. STAs are never associated with a portal. An AP is not a portable. An STA cannot talk to a portal. No STA can communicate with STAs outside of an ad-hoc.
Tom K: if more than 2 people in an ad-hoc, then they have to all be associated all each other.
Dave B: I don't know. If I only want to talk to one of them I may not have to be connected to all of them. Maybe I drop one association before making another. If one STA takes over, then you now have an infrastructure.

Slide 9 - this is the heart of the matter. This is why we don't want to define the implementation but just the services.

John C: we had this discussion a long time ago. We got to the point of saying the DSS operated at or below the MAC and it could be a unique wiring system. This needs to be a proposal submitted to the requirements document.
Dave B: this is aimed at issue 5.1, supportive of the option 'just the things we need from it'. The asking for these services happens in the DSS not inside the MAC.
John C: handling mobility in a cellular system can imply some things about the abilities inside the AP.
Dave B: to do re-association has an implications on the MAC algorithms, but you can't deny that the service is needed.

Slide 10 - summarized that we should specify DSS services only.

Allan: is this acknowledging that the DSS is a higher level protocol responsibility?
Dave B: it is saying that 802.11 doesn't need to specify the implementation. This is irrelevant to us. I just know that to be an 802.11 compliant system it must provide these services to me. The only thing in the 802.11 domain is what are the services. Higher or lower levels, we don't care.
Tom Siep: authentication, privacy - are these really negotiation of these services rather than those services?
Dave B: yes. Sorry.

Lunch break, 12:30 - 1:30

Dave listed specifically the pros and cons his presentation addressed for issues 5.1, 5.2, 5.5, 6.5, 11.2, 12.2, 15.4 .

Tom K: 15.4 you say that mobility is the only thing specific to a WLANs. What about the hidden node problem?
Dave B: you're right that is another.

Also issues 4.1, 4.2 .

Carolyn: about issue 4.2 - Carolyn disagrees that a DSS is an infrastructure. If there is a DSS in an ad-hoc network, then it can't be by definition. There would then be no such thing as an ad-hoc network..
Wim: are DSSs tied to a DS?
Bob C: the DS is in some way very real, but a MAC cannot necessarily tell that you're tied to it. DS functions could be present in an ad-hoc. At least we should talk about it that way.
Dave B: this is a first cut at what DS has to do, more thinking is required.

The Use of Terms for Expressing the Concepts of "Roaming", "Hand-off", "Registration", and "Identification" in WLAN Systems, IEEE P802.11-92/126, by John Corey

This is a group of definitions which are not in the requirements document, but are needed so we can extend our vocabulary with a common definition of terms. John tried to get terms similar to those of cellular telephony. There aren't good terms for transitions between service areas (micro cells, whether they overlap or not), moving between cells and maintaining a seamless connection. Terms for allowing it and for the mechanism for doing it do not exist. The public and private network access problems - identification for authorization, billing information, etc. Looking at the different forms in which this infrastructure can be applied helped to get an idea of the things associated with performing these functions. Used the word administration rather than management to avoid 80.1 confusion - administration over a particular domain.

Registration - buy your telephone, register it and get and id, get a second id if you register with a second network. Once you're known to that network, registered, you will be recognized by that network, and get processed and billed. This is important because we need to make sure that people using a network are known to that network. Dave's authentication is what I call sign-on and identification.

unidentified: implying that in MAN, ad-hoc ?
John C.: no, in infrastructure 802.11 - in a building or on a campus. Only known users get service. You can get authorization by applying for it - i.e. a human registration process like a cellular system.
Dave B: when I buy a phone and it gets registered it has an electronic serial number that is associated with me - I am the user, I get billed. The device is identified and associated with a user.
John C.: when you register your phone you are only specifying the person getting billed. The user is not identified.
Dave B: assumption is whoever is using the phone its OK for me to be billed. The difference between the device and the user of the device may be needed.
John C.: agrees. Logging on to a network will do a user identification, higher function. When I'm finished I can sign-off.

Handoff - this seems universally accepted as meaning making a transition form one area to another. After sign-on I can make this transition active or inactive, but I have been identified and authorized to use the network. Doing that (moving from area to area) in an inactive state, not much happens. Somewhere within the administrative domain some one may want to monitor my location. If I'm active, and want a seamless transition ...

Tim K: what does inactive mean?
John C.: you can be listening, but you're not expecting anything. No outstanding request for services at the higher layers.
Tim K: easier to implement if you only have to register if you want to do something (i.e. rx messages). May want to pass through rooms inactively.
John C.: the hand-off function has to have a way of finding me if something arrives for me unexpectedly. A mobile robot needs to be active all the time
Tom T: 802.11 needs to know that?
John C.: perhaps. If you have a wireless interface but no mobility, perhaps 802.11 wants to know that, so mobility support can have extra functions.
Tom S: MAC may need to know about a session level connection?
John C.: active and involved with something that needs continuity of service, versus idle - We may need to know that. These are things to consider for the definition of hand-off. If we come up with an implementation we may decide we don't need to know this. John is saying that this must be considered in the definition. There may also be an intra-domain handoff.

The proposed term for transiting cell to cell within a network, is ranging - John couldn't find an existing term for that.

Dave B: is it desirable to be able to wander from one network owned by person one to a network owned by another person supplied by 802.11?
John C.: we may have to know about it. John is trying not to define whether implemented within 802.11 or not. Wants to define a standard which allows these transitions. Log on in one network and walk to another on which I am authorized, and I should not have to log on again. Also I should get shut out of if I don't have authorization.
Dave B: ownership boundaries vs. distribution system boundaries - Dave has been assuming an ESS is owned by one entity. You are questioning whether should we make that assumption or whether we should make inter-ESS transitions smooth?
John C.: one ESS is an administrative domain. Someone can have multiple ESSs which are not connected but have transition between them desirable. We have to define how they're implemented, above 802.11 or not. We must address this in the functional requirement document because the are interoperability issues which don't occur in wired LANs.

There is an issue regarding seamlessness which must be addressed in hand-off definition. Hand-off between service areas within an administrative domain is ranging. Roaming is transiting domains. A unit not leaving the service area, but moving that is loitering. There are also stationary units (mostly APs assumably).

If we offer time-bounded services (TBSs), circuit switched services, process control oriented services - some things like hand-off, sign-on and authorization may take on a different meaning. A dumb STA - we don't want it to have to be too complex

The requirements document needs more depth, that is why these terms are offered. John would like to see some discussion on public and private domain terms.

Dave B: private/public - from the point of view of what were doing why is there any distinction? Isn't it just 2 different owners?
John C.: very little distinction on the MAC or PHY. But maybe on things like the NPRM and whether this standard is applicable to a public tarriffed system. Only for interoperability and only if used some day in public WLANs. Nothing to do with MAC and PHY development.
Dave B: where is the LAN-WAN boundary crossed and when does it not concern us?
John C.: as long as it is covered by LAN, even if public, it is our concern. What is public and private? What about a public university?
Dave B: this is why I tend to model something as authentication - inherent in that is identification. I am telling you something I told you is correct or you are deciding to believe me. In both public and private cases this is true - how it is set-up is nothing to do with us. Not necessary for us to deal with authorizing me to use you.
John C.: authorization is making sure the computer is allowed to participate in the co-ordination function. Whether the authorization is for me or for the device is outside our responsibility.
Tim K: we don't need to care. I can hook up to any Ethernet system, the network stops me from connecting.
Dave B: in wired systems I have to beat the physical security to get the connection hook up, but for the WLAN I may even be able to listen promiscuously by driving into the parking lot.
John C.: WLAN must be able to tell a unit stop reporting things you can hear to your user. To gain access you have to have security.
Tim K: you get associated, this is like making an Ethernet plug active. You have access - it does not mean you can do anything yet. (Dave says yes). OK. in WLAN you could be unauthorized and still suck up bandwidth. Jamming.
John C.: in the presence of the infrastructure if someone wants to establish an ad-hoc network - there are arguments for and against this. Does a WLAN owner own the medium?
Dave B: this is why the concept of a co-ordination function (CF). Just because an ad-hoc and an infrastructure network can coexist doesn't mean the ad-hoc is using the infrastructure.
Tom T: roaming and ranging - MAC function or not?
John C.: hand-off is a MAC function.
Tom T: what highest level involved here?

John C.: maybe 802.1, certainly higher than the MAC layer if you want interoperability. There is a line between MAC services and services the MAC needs - like authentication needed from 802.10 .
KS Natarayan: do these definitions overlap?
John C.: they are all new, and they don't overlap with each other.
KS: autonomous network definition is something like your domain definition.
John C.: maybe we need a definition of autonomous network
Dave B: having definitions is very valuable. Putting all these in the functional requirements document may not be right. To get additional functional requirements we close issues. What is your intent - to have us adopt these, or present them for us to thinking about and eventually adopting?
John C.: expect people to spend time considering if these are practical, reasonable or applicable. If they can be used to write functional requirements for service definitions - trying to get enough words so we can talk on next layer definition.
Dave B: think about these subjects (like ownership boundaries vs. hard boundaries) Subject at next meeting, adopting terms after some thought.
Tom S: separate out definitions from the functional requirements, they tend to get confused. Definitions should come first.
Dave B: we could create an issue at the next meeting that says 'what is meant by x', and if we answer it and close the issue that is an official adoption.
François: John has listed related issue numbers - do I have to extract issues out of that text?
Dave B: people who make presentation are responsible for writing up issues arguments and giving them to François.
Wim: a document presented cannot be the argument that goes into the list - some conclusion derived from that can reference the document.
François: if you give me a hard copy and circle the issues arguments that I could do.
unidentified: have we defined the boundary between WAN and LAN.
Dave B: I haven't - for instance, what is the Internet? (there many different answers from around the room).

General Discussion, led by Dave Bagby

There seems to be very little controversy over issue 5.1 - alternative B seems to be the choice. We have had presentations about it and there feels like a consensus.

Straw poll - those in favor of alternative A - 1; of alternative B -13; abstain - 0

unidentified: when you say DSS - does this mean anything at all? All the services we want regardless of what the service would be world-wide.
Dave B: there is a DS - we are trying to decide do we specify the internal workings of that black box or just the way to use it. If a device is compliant it has to have these functions, or it has to be implemented this way.
unidentified: when these services have been specified will the time requirement be specified? Like how quickly hand-off must be done, etc.
Dave B: we haven't specified details yet, haven't decided which functions those are - time might be one.
Tom T: where are these functions coming from? Is someone else supplying them? Shouldn't we define one and say here's an example of one?
Dave B: maybe. The important thing here is the exposed interface. An example can be good, but its an example not a definition.
unidentified: is concerned. Suppose there is a time requirement - if we come out with a black box specification with a timing constraint, for a user to know whether his distribution system meets those requirements or not, is difficult As opposed to specifying a a set of components which define the DSS in such a way you will have the speed required.
Dave B: personally - understands the concern and it is a worry about performance constraints. Why - TBSs is one obvious reason. Previously it has been felt the TBS may be an option because we can facilitate this at the MAC layer, but there is no LLC layer that can invoke that facility to date. Metrics about the DSS may need to be available. It could be that TBS works in the BSS but not over the ESS due

to a slow DSS, but we don't want to demand that the entire existing infrastructure be replaced. Multiple DSSs could be available according to the level of service you need to provide.

To stimulate discussion Dave begins to present thing he feels have been assumptions made by the group to date:

(1) APs are physically stationary.

Tom P: practically, an AP should be stationary with respect to the station. Don't foresee them moving with respect to each other
Tom S: seeming not to be stationary could occur due to an AP failure when a backup takes over the AP could seem to have moves.
Dave B: STA just reassociate due to the infrastructure moving.
François: assuming that the DS can be any existing or future DS, and the DS is also wireless, then AP could move.
Tom T: a moving AP is related to the ad-hoc - if every node has AP functions then they move.
unidentified: why would we want to assume this now. Mobile AP concept may be useful in an ad-hoc situation Or what about nodes being out of ranges and using a node as a range extender. If this is an AP (the range extender). why place the stationary restriction now.
Tom S: someone may be a passive hand-off node without knowing.
François: in ad-hoc example its not an AP function, but more like a cf. Is a CF an AP or vice versa?
Wim: can imagine an AP function residing in a notebook in an ad-hoc network. Then its not an ad-hoc anymore, it could be a wireless AP.
Greg: how does an AP know its mobile. How does it know the difference between the nodes moving and it moving?
Dave B: If relative movement between the two is zero or small, then this just looks like propagation changes. But above that only the relative motion matters.
Tom T: do we care? The relative movement is all that matters.

(2) DS to be built out of whatever you already have.

Wim: does it also provide for TBS with whatever you have? Does it apply to all the services?
Tom S: what should a minimal DS be able to be built out of?
unidentified: this is simply false - we will place requirements on the DS.
Dave B: yes.
Greg: we have said we are going to say you need services to get 802.11 - that's expecting the world to adjust to us [sec note: many people in the room say exactly the contrary - that is us adjusting to the rest of the world.]

(3) inside the DS an address == location

Dave B: in existing networks, address 5 specifies the addressee and the location.
Tom P: location physically or routing?
Dave B: physically not routing.
Chan: the standard shouldn't get involved in address location. It shouldn't be an exposed function.
Dave B: personally I agree. Probably to figure out how to deliver something in a DS you need that. People tend to assume this for existing networks.
Chan: this is true for telecom networks. Not so sure for LANs.
Tom T: define location and address?
Dave B: an Ethernet board has a location and will throughout its in operation , as long as the board is plugged in.
unidentified: given that, its not quite true to say that intra-DS addresses are requires. If APs are connected to the Ethernet as MAC level bridges there is no addressing of the APs required to deliver those MSDUs. It is the addresses being serviced.
Dave B: read the minutes of last meeting - APs and MAC layer bridges are different.

(5) APs are addressable in the DS address space.

Chan: has never seen a need in his models to address in AP except as in internal hardware function. It is true that certain things go from a processor to a port supporting an AP. There is no case where as a result of communication with the outside the address-ability of the AP is accessible.
Dave B: it depends on which side of the boundary you're looking from.
Chan: if a message has gotten into the DS addressed to an STA, the AP - there is a medium between the AP and the portals, so there has it be an id for it to get there. In my view the thing called DS - if you put in a message at the portal and the content is only the LAN address of the destination, logic somewhere in the DS says that destination is a a patroller place and this is never visible outside the DS. Basic condition if an incoming message encounters a point where it must make a choice of destination then it is internal.
Bob C: but the DS is made out of today's LANs.
Chan: I don't except that. Multiple APs don't have to be linked by existing LANs. We have argued about models. One is that taps on a backbone are one APs. I state that that is unworkable for speed and efficiency. The alternative is multiple AP (this is only an antenna) the high speed medium is a backplane and the long piece of wires carries only messages for one AP. Addressing - it is difficult to have an AP per tap on a backbone without giving that AP an identity - that is unliveable.
Tom P: agrees with Chan's opinion - the system is not workable unless the backbone LAN makes fiber channel look slow. To figure out what's happening dynamically you will need a LAN that make 100 Mbit/s look slow - you need a fast connection to the decision makers.
Bob C: the last point - the thing that distinguishes a WLAN from a LAN is that a given address is not attached to a given "bridge". Given that fact the we just said normal LANs will be the DS. There will be a protocol between portal that says what address am I now servicing. Bob disagrees that the bandwidth requirement for this is so bad. We can't impose a protocol on the existing LAN stack that doesn't exist today.
Wim: is lost. Aren't we going back to point 2? Tom and Chan don't agree?
Tom P: Thinks its an admirable goal, but you will find you can't get enough speed. Think how many APs it will take to cover a building - with these APs all talking to each there will be too much added traffic to the existing LAN. Since all messages go everywhere in the LAN, then much information will have to go for hand-off and routing - this will be more than can be supported by the LAN. Chan's idea of tacking existing wiring to a centralized point for making these decision will be required.
Tom T: if the AP is a bridge then no intelligence is required. Registration is done at the radio level. the AP can take packets off the backbone for nodes under it.
Chan: the hand-off function implies there is only one acceptable route at any instance between the STA and the infrastructure. So much redundancy will be required that there will be more than on AP in the room. APs must be used in concert if either receivers receive the message successfully - this cannot be done over a wired LAN.
Carolyn: Thinks there is confusion about just where an AP fits. It may take many antenna to cover an area - this is certainly true with IR where multi-directional coverage is desirable to overcome shadowing. These may all connection back to one intelligent headend - to the 'god almighty'. This is where the AP is - not at each antenna.
Chan: if you have distributed antenna which are not APs it is easy to do with optics. Summing is a small problem. Once you have a radio system combining analog signals from multiple systems - choosing between them on a logical level versus an analog level is required. The logical level is easier for radio.
Dave B: pre-suppose that an STA(1) is associated with an AP and another STA(4) with another AP. The sequence of events when STA(1) wants to sends a message for STA(4). So he transmits. Then what he sends is received by the AP. It is possibly received by other APs too, but only processed by the one it's associated with.
Wim: in Chan's concept an STA is associated with multiple APs.
Chan: there's is a designated AP, but that is used primarily for messages toward the STA. The first attempt at this message may be made at the designated AP. If any other AP gets it successfully and the designate AP doesn't then ...

Dave B: let's assume that there is a one to one mapping between STA and AP for now. the AP receives a message. If it came from an STA which is associated now, the AP gives this thing into the DS. This is straight forward for one to one STA to AP association. Suppose multiple associations - 2 or more APs. It starts the same way - the STA transmits, the APs receive. They each check for association and pass to the DS if they are associated with the STA. In the second case the DS gets multiple copies and needs to figure out what to do. This is just a small detail.
Wim: this could create a lot of difficulty. Wim thinks that the APs need to work out among them who forwards the packet, so that only one of them does.
Dave B: OK, the DS now has the packet, not the APs. The DS now has one or multiple copies. You might get 2 copies of this then, but something got into the DS.
Bob C: this is where the assumption of standard LANs comes into play. Bob supports Wim - if a STA can be heard by more than one AP the APs must sort that out. Standard LANs are not set up to sort out multiple packets.
Jim: the AP needs to make a decision before forwarding to the DS - is STA(4) within its BSS. In some architectures that's not a problem. If all STAs can hear each other then its done. But if it's a store and forward AP then it needs to forward to STA(4).
Dave B: let's just say it has got to the DS - the DS might say it doesn't need to go anywhere, ...
Chan: there may be a qualifier on the STAs data which the AP needs to look into before deciding what to do with respect to the DS.
Dave B: you are missing information from a discussion this morning - Dave quickly re-iterates some of this morning's discussion. I am trying to describe a small sequence of events that got something from an STA into the DS. The DS tries to find the AP to which this should go - if the DS one that sends everything to everybody then its simple, though inefficient.
unidentified: if APs operate in promiscuous mode on that LAN this is efficient enough. That's a possible model for a DS that shouldn't be precluded.
Dave B: once you determine where to send it you send it. Talking from concepts only - don't care how, only saying that getting it there is the action required.
Tom T: interoperability of APs from various vendors could be a problem if the DS is disassociated.
Tom P: have to define what the intra-DS protocol is that is used for the non-handoff. There will be a lot of information. You can always disappear an interface. I understand that the real media interface is the only interface that counts. We can define the structure between the two, but there will be market pressures to have collocated WLANs with different media. These are reasons to standardize the DS. We should worry about the DS after we're sure that market hasn't evaporated while we took too long.
unidentified: are you assuming that after the STA sends a message to the AP there is no ack? Does the AP know that the message got there?
Dave B: I don't care for this discussion. I was just getting global concepts of getting a message from STA(1) to STA(4) over a DS. I was making no statement about local ack versus end to end ack.
Bob C: we're getting hung up on a performance issue that's not as bad as we think it is. If APs recognize 'n' system addresses as addresses they should accept, the amount of traffic they have to accept will be small. But if we have explicit instructions for every frame then the traffic will be too high. But it is outside of the DS - standard LAN will have support for this.
Dave B: how it gets from STA(1) to STA(4) over the DS is outside our scope.
Wim: shouldn't we try to fill it in with examples of the DS with whatever people have available? To check whether what we develop here is feasible?
Dave B: we use the association information to make decisions about what AP is associated with what STA. How the DS does that is irrelevant.
Tom K: if this is true, then single STA associated with multiple APs isn't your business. If you only care that STA(4) doesn't get redundant packet.
Dave B: could make the case that STA(40 would have handle that.
unidentified: to my knowledge no standard IEEE LANs allow duplicates - some higher levels won't support this. It must be required that this does not happen.
Chan: it must be an absolute requirement that by the time packets get to the outside world there are no duplicates.

unidentified: this morning we decided the DS would be existing equipment.
Chan: if by DS you mean the twisted pair wiring I agree. If you say that the DS implies we must work with 10baseT I don't agree.
Dave B: (recaps bit of earlier again for Chan) You could assume that each AP was a Cray and they are connected by simple cable, or simple cable connected to one major brain. It doesn't matter.
Chan: if you imagine that the DS is something like 10baseT then the STA/air interface must not exceed that interface.
Dave B: only if you dump everything on there.
Chan: APs can be made with varying levels of intelligence.
Dave B: returning to the simple getting data from STA(1) to STA(4) story - AP gets the message for STA(4) and checks to see if STA(4) is associated with him. If yes, sends it on the wireless medium. If STA(4) is no longer associated with you, maybe he left, one condition is to give the data back to the DS.
Tom P: another case - if the AP has queuing then you could hold it because STA(4) has been associated with you.
Dave B: the closer to transmission time you look for association the better off you are. This is the basic sequence. Think about what to de next (i.e. as a group).
Wim: is having second thoughts on the vote on issue 5.1
Dave B: I sensed that too, the group doesn't want the to go to the plenary yet with a stand on that issue.

Straw poll: (people who still think the same 12), (people who have changed their minds 2), (abstain 3). Dave can report the vote to the plenary, he was looking for a recommendation, but the absorption rate is slower than that.

Meeting adjourned: 5:15 PM.

Thursday, November 12 1992

Meeting called to order at AM, by the chairman Dave Bagby. Carolyn Heide secretary.

Comments on Document IEEE P802.11-92/98 WLANs Station Management Service by Steve Chen IEEE 802.11-92/124, by François Simon

Steve defined 3 objects that could be managed - MAC, PHY and PMD (PHY media dependent). François has taken those objects and thought about what could be added.

There is frame management, which could recognize frames different from regular data, and manage peer to peer.

There is the MAC management function which would do the normal MAC management, and for WLAN there is probably some function it also needs to do to support authentication and registration. These functions may be higher in the protocol stack, but a portion of it then may have to be handled in the MAC. Security, maybe this isn't the right word, but some function like this may have to be here.

There is physical connection management, since we have different PHYs, which could involve frequency selection and the related issue of interference management.

Co-ordination management provides information to local and external management functions.

We started with a MAC/PHY model, and we have added management to that. This provides more detail - it does not modify that model. Basically this is argument for issue 13.1

François wants to know if it is acceptable to this group to add these arguments to 13.1?

Discussion

Jim: this provides a fantastic framework. Steve brings up the topic of you agree a hopping sequence in an FHA PHY - it doesn't fit in PHY or MAC, so it must be in station management. But there is a catch-22, how do you synchronize station management's if you can't set you hopping pattern yet. Also we don't want to pollute the frame by adding PHY specific information. Maybe a generic control field is the idea.
François: those concerns are farther down the line - can this be the framework to expand upon? The station management may be all we need. François has talked to Steve and there may be some more points be added, but it is contentious where they will reside. The management will have to take care of some of registration and authentication, and a portion of security too.
Tom S: on Steve's diagram, is confused by the functional blocks. On the right side is the management stack, but there isn't a block on that stack that interfaces directly to the frame management - does frame management just it ease drop?
François: this is a logical path not a physical one. An external agent would come in at the frame management. A management agent may want to know things like frame counts, etc. This would come as a management frame not a data one - so the frame management is a filter between data and management frames.
Tom S: it may completely remove frames, or modify them?
François: it will see all frames going by and pass all of them through.
Wim: instead of frame management it is a management frame filter.
François: it is interested in all MPDUs passing by. It does control recognition.
Wim: as to statistics counting - no, it captures management frames. It directs management frames to the co-ordination function.
Bob C: isn't it more like a management SAP?
Johnny: it is a SAP - all station management frames go through it.
Bob C: seems like an address trap more than a function.
François: yes, that's the way I see it.
Wim: So its a management SAP, not a function?
François: yes.

Issue 13.1 - looking for a quick indication if people think this is the right approach to attack this issue. The only alternative so far is the combination of this document and Steve's (92/98 plus 92/124).

straw poll: those for this approach (9, 0, 12). Basically this reflects the fact that people haven't had time to read it.

Wayne: attended the PHY group yesterday, and there is no consensus or agreement. It may be presumptuous to assume we can have one MAC. Abstained because we haven't got our arms around this at all. For example, adaptive power, antenna diversity - special frames are need for these.
François: agrees. In this model many of these points are missing. But the blocks into which these things fit are there.
Wayne: concerned that the framework have enough open tokens so that as we define attributes we have a place to put them.
Jim: that is a functional detail. We need to get the blocks first.
Wayne: we need to work both ways, but we may not converge anywhere near the middle.
François: you made a comment about one MAC not working. That is part of the PAR, we have no choice.
Wayne: it's quite a challenge.
François: feels that all the functions needed would fit in here. We want to agree on an empty tower of objects and fill in the exact functions where they belong.
Wayne: there may be a two way exchange from the PHY. Like a hard coded id number in the PHY that the MAC needs to know. Where does that reside and how does it get passed around?
Tom P: what is the purpose of this number? Does it have anything to do with the address? Or is it the id of the hardware device?
Wayne: could be anything. One theory is a unique id for the terminal.
Chan: is afraid we are ignoring voice/data integration - or connection service type operation. A MUX is needed to somewhere to separate isochronous and asynchronous services.

François: why does this affect management?
Chan: decisions are needed - at what point in the chain do we separate isochronous and asynchronous data. Somewhere there needs to be a MUX. Maybe it is the next block up in the picture.
Dave B: is hearing "I don't want to agree to anything because I don't know all the details yet". So this subject obviously needs more discussion. and thought at a later date.

Dave wants to talk about things to be reported to the plenary. There is the list of papers presented and issues impacted by those. Also at last meeting there were recommendations to close some issues, so this must be done now. These are:

Issue 1.4 - very strong sentiment for being careful about too many issues.

Issue 15.2 - vote to close this as meaningless because the 2 things it asks about are orthogonal.

Bob C: thought we said yesterday that we might be able to provide TBS within a BSS but not across an ESS?
Dave B: the keyword in this issue is "coexist". What does this mean?
Jim: the origin of that phase is the PAR. We decided last meeting it is meaningless.
Chan: it could have meant in a collocated scenario. That is a valid issue and a nasty one, but a different one.

Issue 16.10 - OK

A recommendation that Dave thinks came out of this meeting - we talked about what the DS is built from, and whether to specify implementation or just services. The first we didn't conclude, the latter we had two strong votes saying we think we should specify service definitions only.

Bob C: has a question on issues. Thought that there was good consensus on the paper where he attempted to capture the feeling about the MAC/PHY interface at the Minneapolis meeting, yet there were no issues closed.
Tom P: joint issues like the interface need to be closed in joint MAC/PHY or full working group meeting.

Goals for next meeting. think about :

(1) are there any other categories of services needed in the DSS. i.e. what about TBS services.

Jim: address translation
Wim: what is the "where is the DS boundary".
Jim: that's not a DS service, that's a new heading.

(2) network management functionality should be talked about - what is needed and why. Would like to see some submissions on network management.

(3) authentication and privacy - why are they needed and what are they.

(4) ad-hoc - what does this mean? Infrastructure has been well thought about, but we can't define how these two interact until we know what they both are.

(5) the parametric model of the MAC/PHY interface. The PHY group is driving most of that work, and we need to participate.

Jim: we're going more top down, but we're doing it a bit at a time.
Dave B: but we need to get our words in before they feel its settled.

(6) comparison of MAC proposals versus services. These MACs contain the end results of proposed functionality. We will probably never pick just one outright, but we need to evaluate how each supports the functionality we need, so we can pick pieces from them.

Wim: do you think we are ready for that?

Dave B: no, but he is planting the idea, so people can take the proposals and offer up comparisons.
Wim: under what criteria? do we have enough details of the services yet?
Tim K: we should require in all new MAC proposals that they specify how they provide services so the group doesn't have to it.
Dave B: what does the group think?
Wayne: could that be a revisiting of 92/122, the draft standard outline?
Dave B: that is a subject for this PM in plenary.

It appears that people think its too soon for this MAC analysis, so we'll put it off.

(7) is there any other subject that people want to cover or submission they want to present. Now we are setting the things that will be done next meeting - anything not decided on now may not get time allocated to it.

Bob C: we should take gross cut at the services for TBS, not restricted to DSS. And if we don't close something about the MAC/PHY interface we should address that.

Discussion of issues the MAC/PHY Interface, driven by submission IEEE P802.11-92/100 Proposed WLAN Architecture Bob Crowder. In this paper Bob set down the decisions he believes were made about the MAC/PHY interface model. he believes we may be able to close issues - 12.1, 12.2, 12.3, 12.4

12.3 -alternatives: smart, dumb

Wayne: how dumb is dumb? It has to at least be able to service requests doesn't it?
Bob C: tx, rx, clock is as dumb as you can get.
Dave B: recap - the PHY wants to be completely un-intelligent, so the MAC needs to know how to twiddle knobs. But with one MAC many PHYs we need to know what knobs.
Tom P: why does reading a PHY type out of the PHY mean the PHY is intelligent?
Dave B: it doesn't but it does mean that the MAC has implicit knowledge of the PHY actions required by PHY type.
Tom P: it will take them longer to develop PHYS than it takes you to develop a MAC that can know how to handle anything they develop. In practice a totally dumb PHY probably really won't work though.
Wayne: isn't there a level a little above totally dumb? PHY needs to be able to handle some nature of requests doesn't it? Dumb must at least detect service request type.
Jim: you're assuming a lot here.
Chan: it is very important to have the PHY blind to what service is going through it. It sends and receives messages which are like packets. How those are handled is in the MAC or above. If dumb and this blindness are synonymous then Chan is in favor of dumb.
Wayne: the MAC group seems to think the MAC needs only type and data, Larry thinks power level and such also. Either the intelligence and knowledge of the global system is in the MAC, which requires a smart interface for MAC to get things, or there is a dumb interface and the PHY has to take care of its entire world all by itself.
Chuck: from the viewpoint of the MAC we want the PHY to be a virtual circuit. (Chan says watch you choice of buzz words!). To accomplish that you need flow control - every circuit has different throughput. There are different physical characteristics that the PHY needs to accomplish.
Chan: if we assume you are right, the discussion is not whether there's flow control but whether it passes through the MAC/PHY interface. What we should assume on the MAC/PHY interface is a variety of hooks. The MAC deduces things to tell the PHY from information it gets both from the layer above it and from the PHY..
Wayne: there is an issue in portable, battery powered systems that the PHY people want to have a sleep mode as a least possible function. You want to come awake as early in the data flow as possible. Even a simple any signal detector. A wakeup call would be a broadcast message. So I think there are a little smarts in the PHY that say are required to listen to broadcast.
Tom P: pacing will come from PHY to MAC, just like on conventional CCITT modems. Nothing says that clock coming up from PHY can't stall in a phase that would impose flow control. The concept of

virtual circuit is inappropriate - implies peer to peer connection not inter layer communication. Wake up modes come in very different flavors - if you wake on just activity you will be awake all the time. These problems are better addressed jointly than here.
Jim: was assuming dumb/smart, complex/simple, logical/physical were all matching pairs, but that is not so. Smart is complex and dumb is simple may not be wrong. Dumb interface may be very complex for the PHY. For cost reasons we may not want to do that. Physically would like to see a broader interface separating out some of these functions.
Chan: supports Jim. Parallel commands should go across. Multiplexing clock perhaps, or multiple sleep modes. Just the presence of clock meaning tx one, for example. We must not have in the PHY the obligation to understand the meaning of the bits that pass through it - that's gospel.
Jim: parallelism is very important. You are constrained by time, and there is per packet or maybe even per bit information that needs to go across the interface. The real-time things can be facilitated by parallelism.
Bob C: we are talking about two issues - the width of the MAC/PHY interface and the complexity of the phy.
Chuck: width?
Jim: means pins.
Wim: are we talking about the intelligence level of the interface of the PHY?
Jim: an interface can't have intelligence. We are talking about the intelligence of the PHY. If intelligence can be done in a medium size state machine - that is dumb
Tom P: it must have command recognition to adjust tx power, that is all more than handshake and data passing. If MAC has to segment data from LLC this should be invisible to the PHY. The modem ought to be very simple. The MAC needs signal quality information and presetting of something like gain control and tx power.
Chuck: frequency control needs to be added too.
Chan: that is argumentative.
Wayne: least power consumption should drive the decision. Keeping the MAC awake all the time must be avoided. Certain flags in the preamble could be recognized by the PHY to say this is something wake it up.

12.2 -

Bob C: the diagram in 92/100 shows what we agreed on.
Tim K: the DTE/DCE interface is not the MAC/PHY interface but it is the exposed?
Tom P: standards define interfaces. An implementation exposes what it wishes. An exposure is legal, but you can't force them to expose it, all you can say is if you expose an interface it will look like this.
Bob C: the significant thing about DTE/DCE is that it's an interface between a medium dependent PHY and the independent part. Bob thinks that if there is an exposed interface in the MAC/PHY area this is it.
Tom P: since this is the MAC subgroup. Is it the group's position there is not an exposed interface in the MAC? Bob is proposing we force an exposed interface on the PHY group. Tom thinks that we should say we don't think there's an exposed interface in the MAC, and we suggest to the PHY group that they have on in this area.
Bob C: recommends DTE/DCE interface be defined and exposable and live between the medium independent PHY and the PMDs.
Tom P: agrees.
Jim: agrees.
Straw poll on this recommendation: in favor 15; against 0; abstain 4.

Jim: is there a new issue here? What functions are contained in the medium independent PHY layer?
Dave B: there doesn't seem to be an issue on this now, so François should open a new one.

A brief presentation on model and terminology, by Tom Phinney

This might be more appropriate in a joint group but Tom has sensed a confusion about model and terminology. This is an attempt to get at a formal model and a common vocabulary because he feels people are 95% aligned but cant communicate that.

The first slide attempts to show a building - multistory with a central atrium, walls and pillars and wiring closets shown. Putting radio in this building would require a lot of microcells, overlapping without many clean partitions. Breaking into separate regions is very difficult. You could get four logical pieces tied together with a high speed backbone, perhaps requiring a DS with 4 input points. May have to have inter-floor ties because of the atrium allowing for the potential of inter floor interference. This is to show a real building with a real potential site.

A cluster of DS equipment could be located in the wiring closets with some kind of communication between them. Micro cell overlap and visibility problems Tom calls "coloring" problems. As a model for an STA, Tom thinks of a notebook computer with an antenna. The antenna is logically part of the PHY. That goes to a MAC, and to the LLC and etc. .on up that we don't know and don't care.

AP and DS are the confusing terms.

To suggests that the AP model starts with something that consists of a processing chain - consists of one or more antenna, one PHY layer which recognizes symbols and passes them To a MAC which interprets messages. As soon as you interpret bits you are in the MAC.

For the single STA, there could 4 antenna on the ceiling with which it can communicate. There are 4 processing chains which go to a processing point, a 'bridge', time-space-time switch at the MAC message level. This is not a bridge to the outside world but 4 processing chains that come together at some point - nothing to do with physical implementation, partitioning of thought only. Nothing to do with separation or integration of equipment. A DS consists of a number of processing chains and processes which are not just the switch and bridge functions, but this function can be though of as a traditional bridge and a set of APs. There is association and reassociation between the 4 APs; discarding of data from an AP, and selecting to just discard or reassociate; choosing to accept or reject messages from the AP that is not associated with the STA; there is some duplicate rejection possibility; and alternate path selection.

Microcell usage management is required - how to maximize use and minimize interference. Microcell management is required to choose when to turn which microcells on and off, simultaneously, or in what kind of sequence. This is where the intertime orchestration is required. In each wiring closet there is something that solves the coloring problem. Most of it is solved in a small area -a small piece of this may have to be sent over the wired LAN to another wiring closet but not the majority of it.

No outside gateway to the outside world is shown here. Only when a port to an outside LAN is shown do we get bridge-like functions. This is an enhancement to the basic bridge-like functions. Adaptive power control and inter-AP problems are localized.

Discussion

Wim: in this model does the DS fully include the AP?
Tom P: DS is a collection of APs plus the way they interact.
Wim: the model we have used up to now has an AP which connects to a DS. This model includes the AP in the DS?
Dave B: STA emits, in a conical manner, a signal which may get to 2 antenna, and the towards the STA signal may come from 2 antenna. (Tom interjects that his concept of association is the picking of one of these). Two things came up and somehow they became one coming down. How this choice is done is to be determined. If you constrain your self to one to one its straightforward.
Tom P: one to one in down direction is easy - you can know what antenna to use to get to am STA.
Tom T: what about duplicates - this only works if all APs tied to one headend to do duplicate filtering. How else do the duplicates get sorted out?
Tom P: duplicates must be removed before hitting the existing LANs. Clearly this is a problem
Tom P: so the group has been working under a model where the DS does not include the AP? (group says yes)

Dave B: expressed in my DSS presentation this morning were some people's beliefs that there are performance requirements from the DS. I was trying to say I don't care - defining the information is what's important.
Tom P: to tie many APs to one chunk of DS solves some of these performances worries. The DS can be centralized or distributed. It can be viewed as basic MPDU store and forward functions and interactions within the DS.
Wim: note that the STA model and AP model are identical now. What makes the AP unique? Wouldn't there be a chunk of something that interfaces to the DS in an AP? Is there any DS function in the AP, which makes it different from an STA? An AP would then be a PHY, MAC and something else.
Tom P: probably APs have a DS port.
Bob C: the only issue is how you transport MPDUs. Communication between APs is the same whether they are physically located in the same place or not.
Tom P: we will have to come up with an inter-AP protocol. it may go over a localized DS or out into a existing LAN. But in a building with 450 APs (15 by 15 foot coverage area such as Carolyn says could be expected from an IR antenna), orchestrating the coloring problem is going to be difficult over a LAN.
Dave B: is the thing that connects an AP to a DS a piece of DS that is part of the AP, or is it part of the DS.
Tom P: a DS is not a collection of APs, it is a function that connects them.

Meeting adjourned: at 11:50 AM.