Tenative Minutes of the IEEE 802.11 MAC AdHoc Committee
Raleigh, NC
January 14-16, 1992

Tuesday January 14, 1992

The meeting was called to order by Chairman: Robert Crowder, SHIP STAR Associates. Bill Stevens, Apple Computer, Secretary.

Legend of Names:
Bob: Bob Rosenbaum, WINdata
Carolyn: Carolyn Heide, Spectrix
Chan: Chandos Rypinski, LACE
Chris: T. Chris Baucom, IBM
Crowder: Bob Crowder, Ship Star
Dale: Dale Buchholz, Motorola
Dave: Dave Bagby, Sun Microsystems
Dick: Dick Allen, Apple
Frank: Frank Koperda, IBM
Francois: Francois Simon, IBM
Jonathon: Jonathon Cheah, Hughers Network
John: John Eng
Ken: Ken Biba, Xircom
Nat: K.S. Natarajan, IBM
Randy: Randy Haagens, HP
Simon: Simon Black, Symbionics
Schuessler: James Schuessler, National Semiconductor
Vic: Vic Hayes, NCR
Walt: W.S. Johnson, Phototonics
Wim: Wim Diepstraten, NCR

Ken Biba presented an update to his MAC proposal.

Media independent MAC

No distribution system is required. It is optional, and used for range extension.
Optional ESA distribution system provides:
- MSDU relaying to/from wired LANs and beween adjacent BSAs
- Station "mobility" (Chan pointed out that "roaming" is a loaded word), access control, power control,

Validated performance of proposed protocol via simulation

Requirements:
- IEEE 802 Functional Requirements
- IEEE P802.11 PAR
- **IEEE P802.11 Requirements**

IEEE P802.11 PAR

IEEE P802.11 Requirements

Summary is that there are two basic classes of service required:

- Asynchronous: low avg transfer delay (as low as 2 msec transfer delay)
- Synchronous: low transfer delay variance (MSDU jitter) < = 10%
- IEEE 802.2 LLC support
- Coverage area < 100m or up to 1km

Structure service to match the expected traffic

- Async service for bursty traffic
- Sync service for "realtime" traffic

Structure protocol for extensibility/scalability

- Stations, load, service, coverage

Don't burden all stations with unnecessary capability (there must be some pieces that if it isn't needed, it need not be implemented)

It must effectively deal with the "real" environment. I think it's more important that the data absolutely get through, over absolute performance goals. There is tremendous variation in wireless environment.Protocol must be able to deal effectively with: noise, fades, jammers, overlap,movement.

Increasingly, these products are bought over-the-counter by non-expertcustomers.

Trade optimal capacity/throughput for low avg transfer delay

Support both classes of wireless traffic: async, sync

Provide for : ad hoc, standalone networks
seamless integration into larger networks, etc.

The Three Fold Way

1. The foundation service is a peer-to-peer asynchronous data service requiring no infrastructure, sufficient in itself to support ad hoc networks, yet providng mandatory capabilities for (optional) synchronous data delivery and internetwork service sublayers (2) and (3) (etc....)

Protocol Architecture

- PHY layer
  - Half-duplex peer to peer
  - multiple media
  - single and multichannel PHYs support
- MAC layer
  - async data service sublayer
    - peer to peer
    - augmented LBT with positive - acknowledgement
  - Synchronous data service sublayer, peer to peer
    - reservation TDMA, using asnc service as - mechanism
  - Internetwork Extension Sublayer
    - relaying via wired backbone
    - ~~roaming~~ mobility across wired backbone

access control
Power ~~control~~ management

There is a presumption that between MAC and PHY there is a "uniform" interface so that multiple PHYs may be interfaced to a common MAC.
The LBT approach was previously known as "CSMA", but that term was retired, as it was inappropriate in this context. Chan asks "what do you listen for?" Ken says, "signal - and in some cases frames"

Internetwork extension:

Vic "why do you always say 'wired' when you talk about the backbone?"
Ken "I would allow for other nets. One of the constraints on my system to provide mobility - I expect 802 backbones to provide the interconnect among access points. All I am suggesting is that the "backbone" is "out of band" of the medium used for the WLAN being extended."

The proposal anticipates multiple versions of PHYs - both SS techniques,IR, etc.
Simple interface:

- half duplex
- receive data and clock
- transmit data and clock
- signal detect/channel busy
- channel select
- quality of service

I have run protocols which are degenerate versions of this [protocol], overspread spectrum radios. They do perform in a reasonable manner, though I haven't done an authoritative analysis.
My experience has told me that for direct sequence, we can do tx/rx turn around in about 10 microseconds without great expense.
Most of the time, propagation is good (and there are few hidden stations). In cases of large numbers of hidden stations, performance degrades substantially.
We need to deal with adjacent overlapping BSAs, without common access control.There will be cases of multiple administrations located within overlapping PHYcoverage.
Simplicity, low cost is a high priority.

Approaches:

- LBT core protocol design: nonpersistent with modified binary exponential backoff
- Hidden station enhancements
- Positive MAC acknowledgements
- "hooks" for synchronous service
  - Per-station channel allocation vector prevents asynchronous service transmissions during synchronous allocated time
  - Each MSDU includes a length specifier. Each station accumulates information(adds up the lengths) to develop local knowledge about usage of the medium.Each station also watches for "beacon" frames, which are used to forecast future "sync" traffic allocation.

There are two types of MPDUs.(Ed: This is actually a sequence of MPDUs) (1) Point-to-point MPDU. It uses the "RTS/CTS/MPDU/ACK" handshake. (2). Multicast/Broadcast MPDU. It consists of a simple"RTS/DATA" (one-way). There are no CTS or ACK messages sent (since there are multiple intended receivers of the MPDU).

DISCUSSION:

Wim: How do you recognize that a CTS or ACK is from the intended receiver?
Ken: It is inferred from the timing. There may be a need for a "uniqueness" tag (e.g. sequence number). IT NEEDS further investigation

John: Is scrambling being considered?
Ken: I presume that will be a PHY issue.
John: There may be a need for a scrambler "seed" in the MPDU structure.

Simon: There is a case where two RTS msgs could collide, but causing generation of one CTS (heard by both potential transmitters), thereby causing transmission of two simultaneous MPDUs....
Ken: Yes - I think we do need some way to disambiguate RTS and CTS messages.This needs further study.

Wim: There is an implication that the multicast/broadcast reliability is insufficient to meet the PAR requirements.
Ken: If this is truly an issue, then it is likely that an ACK will be required to meet reliability goals.

Wim: Otherwise we must accept that multicast/broadcast is inherently less reliable than point-to-point
Ken: I think you will find this is true, even on existing wired 802 LANs today.

Dick and Ken discuss that multicast/broadcast, as used today, does not require such high reliability.

Crowder: In my experience, Ken is correct. There are ways of using broadcast/multicast that accomodate their lower reliability.
Frank: there is often a use of multicast/broadcast for group addressing.
Ken: That is indeed what they are.
Mike Cheponis: I'm concerned that the mc/bc doesn't listen for a response from its RTS. It would be better if "sync" stations could in effect say "excuse me- you can't have the channel for that long, because I have synchronous traffic ready".
Ken: In fact, the sync protocol accomodates these needs through several methods.
Chan: Does the beacon have a receive function?
Ken: The beacon is a multicast frame which is emitted by a "beacon station"
Chan: So it is an auxiliary function of an (otherwise) standard station.
Ken: Yes. The protocol uses the asynchronous service for allowing the sync stations to request reserved bandwidth of the beacon station.
Mike Chepponis: Are you restricting beacons (stations) to one per net?

Ken: That can't be done (due to overlapping BSAs). On a single channel system, there must be a way to coordinate the actions of overlapping schedulers. I have begun work on this in my latest paper. I believe there are some simple things we can do in this regard.
Crowder: Is it true that a station cannot send a data frame unless it hears a CTS?
Ken: Yes - a data frame will not be transmitted in the absence of CTS. If the CTS is not heard, either the intended receiver did not hear the RTS, or its CTS was collided, etc.... So data transmission is not attempted.
Carolyn: I'm on the edge of my service area. You're on the edge of your service area. We can talk to each other? There's a conversation going on in my service area, so I can't respond to you (because I see channel busy), and the converse for your BSA?
Ken: This is not the perfectly optimal space re-use plan!
Mike Chepponis: If you set a time limit on how long you wait for a reply, you can limit performance impact.
Ken: This has some of the aspects of a collision detect protocol.
Chan: Is an RTS or CTS a digitally encoded burst (i.e. message)?
Ken: yes. If I see the message length, I'm going to be silent for the whole length.
Ken: You can take an optimistic or pessimistic view of whether the channel is occupied
Dick: You may make the wrong decision, but at worst, a collision is the result.
......discussion of performance implications of RTS/CTS/DATA/ACK (and lack of NAK))
Mike Chepponis: There clearly needs to be a "channel busy" from PHY to MAC which indicates that a valid data signal was received (e.g. an RTS destined foranother station?) in the absence of being able to report a "message" to the MAC.
Ken: Yes - there IS such a channel busy signal in my proposal.
......discussion which leads again to the conclusion that the RTS/CTS/DATA and possibly ACK need to be tied together with some common ID [sequence no., or addresses] to clearly avoid ambiguous situations)
Randy: Just for the consideration of the group. I would separate the "data"PDU from the rest of the sequence, so that if the "data" PDU is heard (and valuable) it can be processed, even if the rest of the sequence is not heard.
Ken: I would need to look at this further, but it isn't something that would make me walk out of the room if it were proposed.
.....discussion of just what is the "length" of the 4-way handshake.
Francois suggests that the "length" fields should be reductive (i.e. each message accounts for the length of itself, plus any other messages that would follow it)
Wim: I would like to go back to the mc/bc issue. The real question we have to know is what kind of applications are really dependent on this feature. This should be included in the functional requirements.
Ken: That is a question that should have been asked (and was not).
Wim: I would like to have a better understanding of what kinds of applications use this.
Ken: (gives some examples) download broadcast. SNMP, source routing, name server
D Allen: RTMP..
Randy: When the medium comes free after a transmission (in a heavy load) many stations may be ready to send an RTS.
Ken: We need a random backoff to avoid this.

Randy: I think a delay after RTS on mc/bc to listen for "quiet" would help avoid collisions (i.e. to make sure its RTS did the intended job).
Randy: If you consider that there's going to be a "pileup " of RTSs, then only the broadcaster will transmit (since we assume he doesn't listen before sending data).
D Allen: You suggested use of non-persistent ....
Ken: Yes - this seems to be the most comfortable.
Randy: So even the RTS will be after a random wait?
Ken: Yes. Performance is quite sensitive (especially at high load) to the backoff algorithm chosen. There's good work to be done in studying this.
Wim: Can you describe what happens when a station addresses a station in another BSA (via the backbone)?
Ken: can I defer this to later, please?

Simulation Results

Peformance

Minimum MPDU Payload	32	(60%)
Maximum MPDU Payload	1500	(40%)
Maximum throughput	86%	

Transfer delay	1.3x normalized MSDU size
@ 10% load	(1.7 msec @ 2 Mb/sec)
Transfer delay	7.5x Normalized MSDU size
@ 50% load	(9.8 msec @ 2 Mbit/sec)

Conclusion: protocol provides 86% throughput, even when offered a load of 4xtotal capacity, with stability.

Buchholz: I'm concerned that your simulations show average performance, but doesn't give information about any specific case. I'm particularly concerned about "starvation".
Nat: This simulation included "overhead", so this (86%) is not actually *usable* throughput?
Ken: That is correct.
Nat: You also did not factor in nonzero BER. What do you think the effect of a (real) environment would have?
Ken: So as "reliability" of the medium exceeds 95%, performance remains good. If it falls below that (sa, 90% or lower) performance degrades quickly. As long as you operate your network below 50% offered load, the (peformance degradation) is not a problem.
Dick: The ACK doesn't actually help the throughput, it just helps the higher layers know (that things worked or not).
Ken: That's correct.
Ken: I came away encouraged by the observation that LBT, in the presence of errors, performs well.

Synchronous MAC Protocol

Goals

- Minimize complexity for asynchronous stations
- Provide low transfer delay variance, allocated bandwidth service

Approach

Defining RTDMA framing structure using asynchronous service MPDUs
Scheduler
Beacon MPDU: synchronizes stations and distributes bandwidth allocation
Per-station bandwidth allocation vector
Synchronous station must implement asynchronous service.
Asynchronous stations must process beacon MSDUs.

Crowder: Is there some sense of a "cycle" by which this synchronous channel allocation is performed?
Ken: Yes. There is a "superframe", starting with a beacon MSDU. During the remainder of the superframe is time which can be allocated for synchronousMSDUs, or contendable for asynchronous MSDUs.
Rosenbaum: You indicated that the beacon MSDU is an asynchronous MSDU? It's error checked?
Ken: Yes - it's a multicast, and it is a standard frame (including CRC).
Rosenbaum: If I miss a beacon, how does this affect system behavior?
Ken: Synchronous activity is sort of like "virtual circuits", and typically changes slowly. So, if a station misses a beacon, but honors the allocation from the previous superframe, things are generally OK.
Rosenbaum: Your simulations did not take into account the "superframes", did they?
Ken: Correct. The simulations are either for asynchronous or synchronous only- no combination of the two have been done at this time. I have estimated the number of (voice) virtual circuits that can be supported. I don't have those numbers, but I concluded that a reasonable number of V.C.s can be supported. There is a "fragmentation" issue, much like memory fragmentation, which affects how efficiently we can pack MSDUs into the available time in a superframe. This is a well-studied problem (wrt memory fragmentation).

MPDU formats:

The Beacon MPDU is merely a multicast MPDU with a special payload field.
The point-to-point MPDU is the same as for async, except for a special case of station-to-scheduler bandwidth request
Generic multicast/broadcast MPDUs - same as async, but with more rigid timing tolerances.

Randy: What if there are *two* beacons?
Ken: We need to figure out how to deal with the intersection of two BSAs with synchronous schedulers. The "asynchronous" service (underlying the syncservice) provides a mechanism for *some* communications among stations in such areas of conflict. So, management techniques could be constructed that could use the async service to coordinate(or at least mutually avoid) multiple synchronous schedulers in areas of overlap (clearly an area for much further study).
Wim: Going back to the case of multiple schedulers which run on a somewhat drifting timebases. There can be some fairly fast drifting between them.
Ken: Stations in the intersection of overlapping BSAs who hear multiple chedulers moving into "time collision" could inform the schedulers that they are drifting together. I haven't given profound thought to this (drift) problem. It seems like it is a soluble problem.
Dick: If your clocks are off by .01% then your clocks will drift about a millisecond in ten seconds.

Simulation of Synchronous Protocol:

Performance:

Minimumm MSDU Payload	32
Maximum MSDU Payload	1500
Maximum Throughput	83%
Transfer Delay	11x Normalizded MSDU size
@ 10% load	(14.3 msec @ 2 Mb/sec)
Transfer Delay	25x Normalized MSDU size
@ 50% load	(32.5 msec @ 2 Mb/sec)

Wim: What's the framing time?
Ken: I think it was on the order of 40 msec to 100 msec.
Wim: Why such a difference between 10% and 50% offered load, since bandwidth is essentially guaranteed?
Ken: The placement of "bandwidth" within the superframe results in large delay (wrt low offered load) because the delay between the start of the super frame and the allocated time grows linearly with offered load (for later-joining stations).
Crowder: Given that the framing rate is sufficiently high, the delay is not really that bad. In fact, it's no worse than existing (wired) systems.

Internetwork Extensions:

(diagram showing spanning tree architecture)
Multiple access points are serviced off of multiple LANs, which then have a backbone architecture as see in wired LANs today.
BSAs are implicitly defined by PHY coverage areas
Access points (using single channel PHYs) essentially create a "sea of partially connected coverage areas". It is not cellular, in the strict sense,because access point coverage areas overlap on the same "channel" (no frequency reuse).
An access point can be configured to offer "hierarchical" service or "peer to peer" service. In hierarchical structures, stations communicate directly with access points only (although some discussion suggested "hybrid" structures would be a good idea as well).
Chan: Is there a filtering function at the access point?
Ken: Yes.
Chan: but it's not a bridge?
Ken: It's not an 802.1b conformant bridge!
Ken: One of the constraints I require is that an access point must always be at the "leaf" of the (spanning) tree.
Chan: Supposing the station transmits and is heard by two access points at the same time. What happens?
Ken: When a station powers up, it listens for "announce" MPDUs from access points. If (on the basis of authentication info, channel properties, etc.) It makes a choice on an Access Point and initiates a registration process to the Access Point. The Access Point and station recognize the presence of each other (is this protocolspecified yet?)
Chan: It "quacks" a lot like a bridge!
Dick: Just as long as it doesn't quack like an 802.1b bridge!

Buchholz: It may quack a lot like a bridge, but you couldn't compare it to the 802.1b spec and call it an 802.1b bridge. Second, if you look at the architecture, you may find that including addressing in all messages makes things easier (from bridging concerns as well).
Crowder: Some other standards allow you to choose whether or not to include addressing in all MPDUs, based upon the particulars of the environment in which the LAN is operated. It's feasible to allow the implementor to choose to include or exclude addresses in all frames.
Ken: Addresses could be included in all frames. For the present time, I have taken the "bit conservation" approach.
Chan: When a station moves from the coverage area of one AP to another, a re-registration must occur.
Ken: Yes.
Chan: And that would take some time.
Ken: Yes, but not that much.
Chan: At any one time, there is only one path "through this thing."
Ken: that's correct. You may lose a frame from time to time while passing through coverage areas...
Chan: Life is like that!
Dick: Especially in the radio environment.
Ken: This scheme gives us mobility that is transparent to the higher level protocols.
Frank: Could you explain the mechanism by which a station moves from one BSA into another BSA? How are they going to determine when to switch, without generating multicast packets?
Ken: There's a special multicast MPDU which is issued by () which resets the spanning tree. When a station moves to a point when it can hear two access points, the station is free to choose to register at the second access point at will. It may also find itself in the state where it loses contact with the first access point before hearing another access point.This would be the pathological case, which would result in a temporary loss of service. In any case,it is up to the station to decide which (of multiple) access point to register with.
...... discussion of source routing and it's difficulties in this environment
Ken: For the vast majority of existing LANs, transparent routing (using spanning tree) are the rule, and this (Ken's) architecture will work. Source routing architectures won't work.
Randy: As soon as you get up to the network layer, this approach doesn't work either. (Ken agrees)
Ken: I'm not trying to purport that this is the ultimate solution.
Buchholz: Implicit in your assumption in moving from one AP to another, is the AP providing some indication of its identity and presence?
Ken: Yes - APs periodically multicast an "announce" which identifies them to the stations.

Access Control:

Stations register with APs and there is an authentication protocol - My (Ken's)preference is a public key cryptosystem.
Security and integrity should be provided by 802.10.

Wim: concerned that authentication applies only to ESAs with access points.Ken argues that there's nothing preventing the implementation of authentication services in all stations, thus providing authentication in ad hoc networks as well.

Crowder: One advantage of this system for high security environments, you could require authentication with an access point for all stations.

Power Management:

A simple protocol is proposed whereby a station tells all others "I'm going to sleep for a while - don't talk to me until (time)"

Dick: You could also add transmit power control so as to use the minimum power which will accomplish the communications desired.

Presentation of the 21 Points:

1. Unauthorized network access on throughput:

Bob Crowder initiates a running conversation considering how "space reuse"efficient Ken's protocol actually is. Discussion of 4-frequency reuse vs.9-frequency reuse ensues (though it is recognized that Ken's proposal is for a single channel only proposal).
Frank: I am concerned that there are some serious deficiencies in the architecture, especially with respect to "handoff" and isochronous traffic -(i.e. handing off synchronous traffic from one access point to another is not a small issue).
Ken offers that there are difficult problems remaining for which he doesn't have a full solution set. He welcomes the committee to move forward in this to a consensus position which meets the needs of as many network types aspossible.

Other Evaluation Criteria

SEE KEN'S SLIDES

Wednesday, January 15, 1992 , Chairman Crowder asked all in attendance for comments, changes, etc. on the minutes from the first MAC ad hoc meeting held in Ft. Lauderdale, November 1991. There were no objections or comments. The minutes are accepted as written.

Chan Rypinski presenting

There are many tradeoffs. I'd like to mention some of the key tradeoffs I have in my mind.
Scalability is important. For example, if you pick an arbitrary range, you quadruple the "area" of coverage by doubling the range. It turns out that in addition to that factor, if you halve the range, you will also be able to reduce delay spread effects, allowing higher speeds without equalization. Also, fading effects are reduced. So, halving the range may have the effect of increasing capacity by as much as 8-12 times. It may be more economical to simply reduce the range than to use more elaborate signal processing, etc....

Let's start with a basic view of technology, and then "we'll see" about range extension by use of more exotic technologies. If you use multiple access points, then it follows that "cost sensitivity" for the access points is implied.
I will talk a lot about infrastructure systems, however I have accepted that"autonomous" operation is mandatory. If you start from an "autonomous" system design, then you will find that you have to add in many extra details in order to design an acceptable"infrastructure" system. This process is therefore additive to the complexity of the "station" design.
The communications between where information is needed and where it is acquired should be minimized. A centralized system achieves this minimization.
Next fundamental philosophical point is:
The space in which these communications take place is bounded. It is bounded by available spectrum and bandwidth thereby achievable.
In evaluating these MAC protocols, I do not think in terms of absolute numbers,such as ms of delay. I've tried to take out (in my thinking) things that could be solved by merely doubling the signalling rate.
One of the fundamental questions is - when is it better to centralize logic and when is it better to distribute logic? The use of terms such as "centralized"should not be taken as "gratuitous acceptance" of such a concept. Nobody should put in a central controller because they "like" them,but rather because it is needed.
These are the main inducements for having an infrastructure:
To provide access and accessibility for external communication by the Stations on the LAN.
To increase the radio coverage available to a Station with minimal radio capability.
To provide organized access control that considers and resolves overlapping radio coverage from the access points.

If you are trying to get "n" stations to talk to a central point, then you have"n" paths to deal with. If you try to get "n" stations to talk directly to each other, then the number of paths becomes much larger (n squared?).
I choose to not use the term "base station," as it seems to be a term to which many are sensitive. I think of an access point as a radio at the end of a twisted pair, probably with a small number of I.C.s.
A ESA should be a closed system. If a station moves from the coverage area of one A.P. to another A.P. within the ESA. Such an act should be a "non-event" to everything outside the domain of the ESA.

There's a further reason for a common hub controller. Knowledge of when a station or access point can transmit is NOT an independent factor for each station. It depends upon the activities in its vicinity. Overlapping is a"nuisance" for many proposers. I attempt to use overlap as "constructive redundancy." It's really very important not "what" was done, but "why" it was done. The degree of centralization is not absolute - it is quite arguable. Many of us see the factors differently.
.....display of a schematic of a large system, with backbone hub, bridges, and access points
For each access point, there is a "bridge-like" device. There may or may not be any active electronics there. It may be just a wiring junction point.Buildings are wired in star-fashion, not in "bus" fashion. There will be a pair of wires between each access point and the "bridge" point.
II show access points (in a sea of stations), coming back to a central "hub controller".

Marvin: I don't think either of those two diagrams relate directly to any of the MACs proposed. It looks like an architecture to me (Carolyn agreed). I don't see how you can say that this architecture can't be implemented using Ken's proposal, or any other for that matter.
Chan: I think you'll have to hear me to the end in order to understand the significance of this to my proposal.
Carolyn: I think the picture you drew matches the picture Ken drew. I don't think either MAC implementation is precluded.
Chan: My proposal is aimed at an economical implementation. It does not imply that other implementations (e.g. bridging onto an 802.3 backbone) won't work.
I'd rather stay out of what other people can't handle and stick to what my system *can* handle.
Rosenbaum: Are you saying that the UTP between APs and hubs a 10BASET link?
Chan: I haven't *said* what goes between the UTP and the APs, but I will.I will assert that the UTP can signal at 16 megabits (because we did that in 802.9).
I should say that what I am getting at is intended to avoid sensitivity to different PHYs. It should work over spread spectrum, narrowband and IR as well.

A Fundamental point:
A lot of systems which provide isochronous systems (and others as well) depend upon the concept of the "time slot."
Coming out of 802.9, where we dealt a lot with IVD, something became apparent to me and that is that there are some serious problems with timeslots. The position of the time slot is a form of an address. That addressing must be negotiated between sender and receiver. This isparticularly difficult in multipoint environments. Thinking long and hard, I think that the best way to do isochronous is the 802 way. The 802 way is that the "time position" of a message is not a form of addressing. Addressing is performed by a header on the message.
If we are talking about multiplexing isochronous services, we are going to do it at a "packet" level.
It will be found that there is quite a similarity between my proposal and all the others. The contents of the message header are derived on common requirements which all MAC proposals will find it necessary to address.

Objectives:
......Detailed on the slides presented-
It should be medium independent. There is no discrimination against optics,for instance. That won't be the same everywhere (else).

The whole thing works by the use of intricate messages. I won't go into the contents of these messages. They would have to be studied in detail. With about 7 or 8 basic messages, you can do almost everything that is needed. For autonomous operation, you can take the basic messages and "tweak" them a bit.
I think we are coming to agreement that a two-way handshake is needed to ensure that the path exists between sender and receiver, and an acknowledgement to ensure the transmission succeeded. I think this is a good sign ofprogress.
There is a "segment" and a "packet" data frame. There is no way we can occupy the wireless medium for long enough to traffic the largest allowable MSDU. So,segmenting is called for.
,,,,,a discussion of Request/Grant vs. RTS/CTS terminology
I needed to know for my purposes "how long is a message", so that I could estimate the performance cost of each message type.
Note: This is a "sequential use" system, where each access point operates "one at a time". They do not overlap in time.
Interfering stations can be put on ten channels. This is like putting them on ten "time slots". I don't mean they are "periodic", just that they occur separated in time.
A cluster is that number of access points that interfere with each other when they transmit at the same time.
The primary mode of averting interference is by sequential use. So, in an installation like a factory, you could have as much as a fourfold redundancy in coverage (due to overlapping access points).
Wim: Can you give us some picture of the timing characteristics?
Chan: If you had ten 1 megabit systems operating simultaneously, you would have an aggregate 10 megabits capacity. If you had one 10 megabit system operating, it would have the same capacity. But the 1 megabit systems would have significantly higher range.
A time serial system could balance the use of the capacity more evenly among those needing it as opposed to the alternative (of multiple [lower rate]simultaneous transmitters).

Central access management and asynchronous access

Dave Bagby takes exception with Chan's repeated assertions that there are somethings that can *only* be done in a central controller. Chan offers -"it's the only way it can be done in *my* system." Dave accepted this as reasonable.)
Bob Crowder: We really need to see *how* the system operates to understand all of the tradeoffs, attributes, etc.
Chan: Any station's messages may be heard by more than one access point. If the access point for which the message was intended receives it without error, then that's all that's needed. If not, the central controller may be able to recover the message from another access point which received it correctly.[this is the constructive redundancy].

Access point Address and Moving Stations

The station listens - it's just now turned on. Is there infrastructure and access points. If so, is this the infrastructure I'm looking for?

REGISTRATION FUNCTION

All through this, there's a discipline. Stations rarely transmit spontaneously. Usually they transmit upon request from the infrastructure.
The INVITATION TO REGISTER message invites the station to register with the infrastructure.

POLLING FUNCTION

There is a calculation (assuming a 10Mbit signalling rate), which in this particular case, results in a poll every 3 msec.
Station failures are particularly difficult to deal with (i.e. just not hearing from a station.) Polling helps the system to keep track of the state of health of the stations.
The frequency of polling is configurable.

Inward REQUEST-GRANT procedure.

The station has a problem in that it must get permission to transmit.Therefore, the situation is different at the hub controller (as it knows much more about what is going on). In the reverse direction, the hub may just"send" without requiring a handshake with the station.
So, while it's necessary for a station to gain permission before transmiting,the converse is not true. It is Chan's understanding that in 802 systems, stations do not have the ability to "refuse" delivery of traffic.
Crowder points out that LLC3 was invented for the purpose of dealing with media with high error rates.
Chan: That's a pretty important matter as to whether or not my protocol is over compensating for the unreliability of the media. If the group feels that a handshake is necessary to ensure reliable operation, I would not object to having it added.
Bagby: The hub only knows about the station immediately following the poll.If any significant time passes between then and when such knowledge is needed,then the knowledge is of little value. In the meantime, the repetitive polling serves to increase battery drain in portables.
Chan: Would you care if the power cost was say, on the order of 1% of the idle drain?
Bagby: I don't think I would mind that, however, having decided that the poll is of no value to the system, I would object to that additional drain. Can you handle a station going away, and shortly there after-registering?
Chan: Yes.
Bagby: Then why don't you dispense with polling and just use the registration procedure instead?
Wim: What does the information obtained from a poll do for the operation of the system?
Chan: It informs the system as to the presence of the station and its relative"strength" as seen by multiple access points.
Dick:Doesn't this violate your premise of "only bits/clock in and out" simplicity in the radio (since RSSI is needed in the receiver)?
Chan: Well, that *is* a problem that needs to be looked at, but at least it's isolated to the *infrastructure*.

Summary of Access Method for Station Originated Packets

Invitation to request is heard from the station's assigned access point.
Request (from station). This involves contention.
A stronger signal can mask a weaker signal.
The weaker signal may in fact be received at another access point (where it is stronger) and queued.
Grant (from access point)
Data (is sent)
Crowder: Determinism is very important. Indeterminate network behavior is even more important to avoid than low performance.
Chan: On radio, NOTHING is certain. Now then, if things are uncertain, you want to sort-of balance things. I think the probability of unresolved contention in this system is quite low.
The sequence at AP is :
INVITATION TO REQUEST
REQUEST

GRANT
DATA
ACK/NAK
Now, direct peer-to-peer:
Station sends message.
Access point delays ACK long enough to hear the destination station ACK (if it does so). If access point hears destination's ACK, then the message is not forwarded through the infrastructure.
The hub doesn't worry about contention because it's in control (no contention).There's a possibility that the message won't be received, but this may be confirmed by lack of ACK from station. The station may have also gone offline since the last poll, which could also cause the message to be dropped. I haven't worked out this case yet.
There's no attempt to send messages to stations which are not registered.
segmentation and auto-grant
There has to be a guaranteed worst case access delay. Without that, it's not possible to say that isochronous traffic will get through. I work almost exclusively with worst case timings, because this is what will occur at peak loads. If the system works at peak, it will work at all lower loadings.
The following areas were presented:
auto-grant
Compatibility and capacity allocation for packet and connection-type traffic.
Excess demand from stations
What happens to frustrated users? For packets, I think you always have to specify where the data "backs up" when demand exceeds throughput. In this case, it occurs in the memory of the station equipment.
autonomous groups not using infrastructure
Two methods:
First station opens up to "look" like an access point. It only needs to issue "invitation to transmit" and "grant" messages. There is no reason to expect that station to actually have logic to deny grants. (a degenerate case)
Simply have a contention-based system. There's one point where I (Chan)have a difference of option with LBT approach. I don't believe it is necessary to"listen" before talking. Sometimes the information obtained is wrong. In all cases, though, the LBT handshake does consume bandwidth.
Simon: some studies have indicated that LBT performs better than ALOHA [which is what the committee has concluded is the nature of Chan's second proposal for autonomous]
......a prolonged discussion as to whether or not Chan's system exhibitsALOHA behavior at the intersecting fringes of adjacent uncoordinated systems.Chan asserts that as long as the loading of the systems is low, this is not a problem.
Simon: in defense of Chan's proposal , the interaction between closely located uncoordinated systems will occur with *any* MAC proposal.Some capacity sharing will occur in all cases, irrespective of the choice of a MAC protocol.
Dick challenges Chan's assertion that uncoordinated closely located LAN interaction is not a problem, while his own design "shares" bandwidth among closely located access points (by staggering transmissions from adjacent accesspoints) is inconsistent. If closely located, uncoordinated works fine, then you could get higher aggregate bandwidth from the system as a whole by not applying the "staggering" rule to access points within the same system (i.e. multiply aggregate bandwidth, not"share" it).
Chan: Isochronous. At the source end, you take the "samples", collect them into groups (say, 48 octets). Once it's accumulated, it should be sent as soon as possible. The only absolute requirement is that it be

sent immediately - just that, it get sent before the next group is accumulated. At the receiver, there is a FIFO which accepts these groups, and feeds them into the destination (device) at a fixed rate.This implies that the receiver must "create" a fixed delay (the queueing delay for packetizing and depacketizing).

This concluded Chan's presentation

Chair asks "What are our next steps?"

** agreed that Jonathan Cheah and K.S. Natarajan are expected to make a presentation in Irvine (in March).

**Does anyone know of any other proposals that we can expect to be forthcoming?(several "possible" proposals were eluded to)

Other suggestions on how to proceed?

Dick: Sooner or later we need to rank proposals on the 21 points.

Crowder: I would suggest we continue to work on ALL the propdsals, with any contributions. The authors of the proposals would fold in contributions from others into their respective proposals. We will not be allowed (by D. Loughry) to pick one proposal until we have a hard, firm agreement on a set of functional requirements. With intense discussion of all the proposals, we may intuitively conclude that a subset of the total proposals were receiving more attention, so that they would be a better place to invest our energies. This is just a suggestion.

Dick: You might also submit suggestions to each of the proposers on how to improve their proposals, thereby constructively contributing to all of the proposals.

Crowder: I still don't have any problem with having Simon present on Jonathan's proposal.

Multiple people feel that agreeing to the functional requirements document is going to consume more time than allocated, thereby eating into time allocated for the MAC subgroup.

What about sequential relationships to the PHY?

Since WINdata has a contribution to present, which may be as much a MAC proposal as a PHY Bob Rosenbaum was "fetched" from the other subgroup, and if he wanted to present to the MAC subgroup tomorrow? He responded, " No. It was intended as a TOTAL contribution, not just MAC or PHY."

Crowder: The PHY group is asking for a specific "sequential use" model of the channel, and questions about how probabilities vary with sequential use of the channel.

Dick: Are you talking about coherence time?

Crowder: Yes.

Wim: I think it's interesting to think about what kind of performance figures we would like to see (from the PHY group). Also, how do we compare different systems' performance. We could start by defining some traffic models.

Dick: We definitely need that. Sooner or later we need to do simulations, and will need those. We may have that already.

Simon: Traffic modeling has a major impact on how things turn out.

Crowder: I bet that not only will there be new proposals, but there will be suggestions for changes to existing proposals. I think we should ask for all of our allotted time tomorrow (PM), and work on these action items we have discussed (e.g. traffic models, ...) We could allocate an hour to get the questions from the PHY group, and a start on traffic modeling...

Simon: Not just traffic modeling. We need to discuss how we can establish a process for "exercising" these proposals in a consistent fashion, so the work can be spread out beyond the individual efforts of the proposal authors.

Agreed: We will work on questions for PHY and Simon's last point.

Meeting adjourned at 5:10pm.

Thursday., January 16, 1992.

Dave Bagby chairing on behalf of Bob Crowder.

Dave: What do we have to do today?

1. Work on questions for PHY group

2. Vic has asked me to discuss our goals for next meeting.

Chan: (discussing management functions)

Scrambling (and block coding). Also loopback testing.

All of these are eventual PHY issues. But it will turn out that these will be "heavy" discussions, i.e. whether these are part of MAC or PHY, but these are part of *something*, for sure!

Also, what do you scramble?

Dave Bagby - what exactly are we supposed to be asking the PHY group in these questions - or did the PHY group have questions to ask us?

Chan: Let's make our own list of questions and issues, regardless of what is "expected" from the PHY. The object is to get the issues out in the open as early as possible.

The first group is "radio control matters."

- Transmitter controls
  - Transmitter on/off
  - Power level
  - Clock and data
- Receiver functions
  - Clock and data
  - Quality of Signal (analog indications)[1]
  - Quality of Signal (bit-stream derived)
  - Channel acquisition
- Management functions
  - In-band/out-of-band "management channel"
  - Scrambling (this is NOT encryption)
    - (clock recovery)
    - (done to avoid spectral "lines")
  - Encryption
  - Channelization (multifrequency) selection
  - Remote loopback for stations
  - Local loopback for access points
  - Jabber control
  - Intelligence Split (between MAC and PHY)
  - PHY I.D. (i.e. which [of multiple possible] PHYs is in use by this station?)
  - Sleep mode

1 Wim asked about "carrier sense." Chan classifies this as an analog QOSindication.

Antenna controls (e.g. diversity, directionality)
Forward Error Correction

Dave: Let's discuss some of these this morning. The two that strike me as more important is "encryption", and "intelligence split". On security, if you rely on 802.10, it takes care of the "payload" security. But authentication probably needs to be handled at a lower level.

Dick: The MAC, we know, is going to have to have a lot of intelligence anyway, so it's more likely that we will want to put such things in the MAC, rather than the PHY.

John Eng: I think encryption should be done at the MAC layer, but scrambling should be done at the PHY layer.

Dick: At Hilton Head, didn't they say that some of the XORs and Shift Registers were in the Mac?

Simon: That would be in common with encryption in other systems (examples), where it's done in the MAC.

Dick: Where is this done in the WaveLAN?

Dave: It's done in a "chip" :-). and in the U.S. :-) I'm trying to get a general idea as to where we think the intelligence should reside.

Jim Sch: I agree with Dick, we shouldn't define functionality that would require a processor in the PHY (for cost and complexity reasons).

Dave: As devil's advocate, - depending on the modulation scheme being used, you may have a processor required in the PHY anyway (i.e.direct sequence). If you're using a simpler PHY (e.g. frequency hopping), this may not be true.

Dick: there may be a counter argument that, at higher speeds, you may argue that you can't depend on a very intelligent interface,where "messages" would be too costly for time reasons.

John Eng: The PHY could be very complex (e.g. DSP), but we should focus on a "clean interface".

Dave : Let me repeat a discussion I had with Dale yesterday. I never imagined the need for a MAC/PHY interface to control antenna selection. Dale?

Dale: It depends on how you do the MAC/PHY interface. If you do it as a simple bit stream, and the PHY simply sends and receives bits, then no such interface is required. But if the PHY is handling"directionality" issues, the PHY may need to look at the destination address for "aiming" the transmitter.

Dave : Is there any disagreement that we should aim for a media independent MAC/PHY interface? If we strive for media independence, then PHY I.D.becomes moot. If we include PHY I.D., then I'm uncomfortable because that would imply that every time we invent a new PHY, it may impact the MAC design. Maybe a key question we should ask them is "do they want a medium dependent, or medium independent - interface?"

Simon: (draws a graph) proposes a "lower layer management entity" which spans PHY, MAC, and slightly above (not specific in how high).

Rosenbaum: Traditionally the management layer has not been used for "realtime" type control issues.

Schuessler: My belief is that this allows the MAC/PHYinterface to be clean, simple, and universal. Then the management entity includes the "complexity".
Walt: Are you then saying that each PHY would have a PHY-specific LLM?
Dave: Perhaps we are just playing a semantic game. Is the LLM really any different than (jigsaw puzzle graphic) a MAC/PHY interface where the "simple" ones just connect directly, and the more complex ones "connect" at more points?
Francois: The management "tower" has been around in virtually all OSI approaches. Unfortunately, it has typically been used as a"garbage can", where things that can't be fit cleanly into the layer interfaces gets "dumped" into the mgmt tower.
Chan: Where we draw the boxes does nothing to lighten the load on what we have to do. When we draw the boxes, we're just setting down a method by which we can explain it. Antenna selection in my experience is a "geographic location" selector, which is a matter not seen outside (what). Once you say that, then the table of "addresses" may live (imaginably) in the management stack, not in the MAC, so we *do* need the box on the side.
Francois: In many cases, the management tower is used to communicate information from one layer to another which is far removed (i.e. not adjacent)to the layer providing the information. This information is typically accessed by an "application" at a much higher layer.
Dave : When I think of information going from MAC to PHY,there's data which goes along with the message - i.e. "this is some data I want to get to 'joe'. If I want to get it to 'joe', there's some 'direction' information which may be needed to aim the information at 'joe'. If the station is mobile, how can anyone *except* the PHY have any information as to which direction to aim to reach 'joe'?"
Chan: It's not just "direction". In some cases, 'joe' has moved to another access point as well. So where does the logic reside to resolve location?This has to be known "in advance" of when you need to send the information. If you have to determine direction/location at the time of transmission, there's a lot of time that will be lost in determining this.
Dave : Geographic vs. logical addressing. What a MAC usually has is "logical" addressing. For the MAC to have "geographical" information is a big mistake (i.e. IP addressing). I think it's crucial to separate geographical and logical addressing in this committee. Otherwise, we're going to be in big trouble.
Chan: If you have that kind of "diversity" (not infrastructure related), then it's not an "addressing" matter. There's a number of things you can do with antenna diversity, which are strictly a PHY matter. Each kind of PHY can take care of this in its own right.
Walt: Could we finesse this by suggesting that transmitters are omni, and receivers use directionality?
Dave : Some transmitters might have omni antennae, and so might some receivers. If a receiver has a directional antenna, and knows how to use it -

great! My gut level is that that's a "dirty issue" that I would like to see the PHY take care of.
Walt: By putting it on the receiver, it's simply a "signal quality"factor.
Dave : I could envision an argument that having it at both ends could be advantageous, and I wouldn't want to give that up .
Chan: I can cite a case (nationwide trucking comms via satellite). It would have been much simpler to use an omni antenna, but they just didn't have the "decibels." Having looked at that, I don't believe many portable computers would want to have that!
Dave : Let's see if we can limit to "what services do we want from the PHY?" If I start with that, then maybe they will say what we need and they will come back with "but we can't." But this would be a start.
John Eng: (makes a case for a "core" set, and everything else is technology and implementation dependent)
Simon: So we would be talking about the minimum set of service primitives.

Service Primitives
Common denominators:

Data
PHY activation/deactivation (i.e. on/off)
PHY status

Simon: argues that PHY status is not a "common denominator". Some PHYs may need to provide this, but not all.
Dave : I would try to think of the MAC as something I can do in software, while the PHY is something that would be implemented in hardware. If there are strict timing requirements at the MAC/PHY interface, then this makes my model very difficult.
.....discussion of whether the preamble is provided by the MAC or the PHY.
.....Differing opinions
Chan: The PHY activation delay (from MAC service request to PHY accepting data) will almost certainly be fixed, in which case the PHY's FIFO will probably be a shift register of fixed length. I believe that eventually we will get to the point where this delay is on the order of one or two bits anyway.
Dave : We need to talk about what we're going to do at MAC WG sessions at next (March) meeting. We also need to talk about simulation methodologies.
Consensus is to address these two issues, and suspend the current discussion. We will talk about "simulation" first.

Dave B turns the floor over to Simon . Also, Ken Biba is invited to contribute.
Simon: A number of people are starting work on MAC simulations or thinking about it. It would be nice to get some idea of, not only some practical modeling parameters we might use, but also discuss tools that could be used, and how we could develop a PHY model to use with our MAC tools.
Ken: Some of the tools we have been looking at allow us to "conclude" the behavior of the PHY, but doing MAC and PHY both dynamically is

challenging. I've been pleased with Extend (on the Macintosh). It runs on the Macintosh, is available for $400,depending on where you buy it. Authored by Imagine That! in San Jose, CA. It is nice because it takes advantage of the Mac user interface. It's essentially GPSS with a graphic interface. It has the ability to include blocks from a substantial library. You can also customize library blocks or create your own blocks. Current version does not use the FPU, but the new version is supposed to use it.

Dave : Are there other tools that have been used?

Simon: Comdisco. A good tool, but it costs $35,000 for the software and Sun workstation.

Wim: Bones (?) Wim has started making his own models.

Ken: I'm intrigued with the possibility of choosing a low-cost common platform upon which we can collectively create a library, where we can exchange parts of our models and build on each others' work.

Wim: What kind of libraries does Extend have?

Ken: It has a variety of libraries, not all of equal value. Some libraries are of exceptional quality, while some are abysmal. A rather intriguing collection.

Simon: We certainly have a number of models available which we would be happy to share.

Ken: I discovered that I couldn't build the model using the standard blocks, or if I did, it would be horrendously slow. The "critical" blocks needed to be hand-crafted.

Dave : Well, that seems to tell people what kind of tools have been used, but I don't think we can make a single choice, which would require people to go out and buy more computers.

Wim: But apart from the tools, I think we could start working on "traffic models", etc. to share in common. Then, everyone could use the tools of choice.

Dave : I suspect that what we need to start is the "traffic models", and these would be related to the specific application areas- so that would be the way to collect the data.

Ken: We could create for each area what the traffic model should be, say, with a histogram. It seems that it should be possible to write that down on a piece of paper. I think it's a wonderful opportunity for an ad hoc subgroup to form(and I'm not volunteering!). One of the key areas to focus on would be "representative environments"(e.g. twin towers). Then we could think of "connectivity tables". Coming up with some common definitions of these common environments would be a strong basis for common simulation activities. This is as important as the traffic models. This could be represented as a "matrix of connectivity".

Dave : Doesn't that also say that you need models for what your distribution system is?

Ken: It depends on how accurate your traffic models will be. One simplification would be to associate individual stations with individual access points, thereby generating some simplifying "parallelism" in the model. I found the most important issue to be the interactions between overlapping coverage areas. Ken (draws graphics) What I've constructed is essentially a matrix with, say, five nodes. And then representing what the connectivity matrix is by recording the probability of one station being able to talk to each other station in a "row" or "column" of the matrix. Then we have a "each station"to "each other station" matrix of connectivity probabilities. This is clearly a "static" representation of one possible environment. You can then create any number of such environments to drive the simulation.

Wim: (draws graphic) From contribution 91-125. I made an absolute model, as much as possible. I just put in an area, I defined a network with a server somewhere. I can put a number of workstations, up to seven per network. For practical purposes, I use two networks with up to seven workstations with one server. I place them on a map, and then I make a table out of that with calculating the nominal attenuation of each path, from any one station to any other station.Now, I have specified the other network to be the same, except I can place the other network by translating along the X and Y axes (shift the origin in the plane). I can also place the other network on another floor (and add say 20db attenuation for the ceiling/floor). Then, to evaluate interference, I can shift the networks together or apart by say 10 - 50 meters at a time. To access the network, I calculate path loss and use a normal distribution on top of that (uncertainty). And I use this to calculate "carrier sense." Once I've done that, I evaluate (at the intended receiver) the attenuation of the signal, and then I look (around the intended receiver) for interfering signals. The signal/interference ratios at the intended receiver (including capture effect)is used to calculate the probability of a successful transmission.

Chan: Do you assume the signal level as constant for the duration of the packet?

Wim: Yes, so the coherence time of the model will not significantly effect the results. But I do evaluate the "produced interference" at the intended receiver by signals which may begin randomly during the transmission being evaluated. So this allows simulation of "jamming" as it would occur in a network with hidden stations. I could add an additional step to include "shadowing".

Dick A: Do you set a fixed attenuation between each pair of stations?

Wim: I recalculate the attenuation each time, applying a random variation based on the standard deviation. The traffic model I used is the Novell network model based on a high load test. For the access, I model our current CSMA/CA (WaveLAN) protocol.

Chan: But you didn't include a "handshake" in the model?

Wim: No.

Chan: Since we're talking about simulation, I'd like to show some modeling .... that was done by the PHY group. Shows one page (on projector) out of one of the references given in the requirements document, which was supplied by the PHY group. This particular reference was concerned with including link-layer functions, such as acknowledgement in the protocol. What I want to talk about, however, is (underlined areas).

Traffic-related assumptions are: 1) heavy traffic, 2)uniform traffic pattern, 3) fixed packet lengths, and 4) fixed synchronization preamble lengths. The heavy-traffic assumption allows queueing effects and user actions to be ignored, since each node will always have at least one packet ready to be transmitted.

I bring up this reference to bring out that this assumptionis noted here. So, for instance, the Novell model may be "heavy traffic", but it may or may not be "heavy traffic" with respect to the particular PHY model being considered.

Wim: For high load, I use a model where all stations are offering load, and traffic is random from any station to any other station. If you use a conjugate server-workstation model, it drastically changes the outcome of the simulation (as the server becomes a significant bottleneck to the overall network activity).

Chan: I would like to point out that there is a potential significant philosophical difference. If the model assumes that all stations can generate an infinite amount of traffic, which is instantaneously available, and this is not necessarily realistic.

March Meeting Plans

We need to allocate time for Nat and Jonathan to present.

Dave : I think we need to have a mind-shift. I would like to encourage us to create submissions to the group outlining what problems we see with the proposals, and suggest improvements (so as to be constructive). This could be much more effective than live on-the-air critique. We (the group) seem to have consensus on:

a.) a short review,
b.) a rundown of the 21 point evaluation , and
c.)then an interactive discussion.

How much time do we need to accomplish this?

For each proposal, there's probably a half day of work.

I would prefer we weight our time toward new proposals which haven't been aired previously.

Remember that the March meeting is likely to be more heavily attended, with many new faces. This is likely to slow down vis-a-vis our (interim meeting) experience.

Bob Crowder (Chairman) said (at a previous meeting) that we would only entertain presentation of proposals that had already been submitted in writing. We wouldn't tolerate "on the fly" enhancement of proposals.

We certainly should allow improvements,but it should be in the form of papers that address, "here's what I changed,and how it improves the proposal."

* Dave Bagley adds his company's proposal to the list.
* Bob Rosenbaum: I'm going to try hard to get our company's proposal in for the March meeting.

Dave : If people would write down (between now and March) the issues that have come to mind during the proposal presentations, it would help in recording the issues and enhance forward progress. I'm concerned that we're not capturing this at present.
Wim: You mean issues with the proposals?
Dave : I think that we should capture issues with or without respect to the proposals. As issues come to mind, they should be captured.
Carolyn: Could we generate a list of documents which are most pertinent to be reviewed before the next meeting?
Wim: Will it be possible to get the new (SUN, Windata)proposals in time to review them before they are presented?
Dave and Rosenbaum: We'll try, but it's not highly probable.

Topics we need to deal with next meeting:

1. We definitely need to spend more time on MAC/PHY interface.
2. Maybe we should encompass "architecture" (as we talked about this morning) in our discussions - not a formal architecture effort, but just spend some time, since it's a topic we will keep coming back to.

***** assume 2 hours for each proposal (2 existing, 2 new) so that's two half-days there. There's also some miscellaneous procedural stuff (minutes, etc.). Also some report/discussion on simulation activities.

MAC block diagrams (like Larry's graphics presented from the PHY group - PHY template).

Security/Authentication. (esp. since 802.10 liaison could occur atplenary).

This adds up to 3 full days. So we'll report what we need (to Vic), take what we get, and spend the first hour of the first MAC WG meeting to prioritize.

Also, we should anticipate time required to respond to MAC-related comments from the Requirements document letter ballot.

Meeting adjourned at 12:05pm.