MAC Minutes
Monday, September 20, 1993

Meeting convenes at 1:20 PM, Dave Bagby presiding. Bill Stevens recording secretary (Carolyn Heide edited later).

We set an ambitious goal last time, and we need to figure out how to accomplish this (selecting a MAC design from among the several candidates). We also have a number of papers to get through.

Basic plan is to deal with the papers as usual. Also we will break into small groups as last meeting. Then, Dave will write decisions and see if we can get through them.

The current candidates:

1. BFP (Bob Crowder)
2. BLAMA (Hitachi)
3. Chan's LAN (Rypinski)
4. CODIAC (Spectrix)
5. Combo of IBM/Xircom (National)
6. GRAP (KC)
7. IBM protocol (IBM)
8. INRIA
9. Slotted Aloha DAMA (J Cheah)
10. WHAT (Xircom)
11. WMAC (NCR/Symbol)
12. Combo of IBM/Xircom (National)

I want to try a little experiment here. Let's try some "anonymous" data gathering. Take your hotel notepad and answer the following questions:

1st question: If you had to pick a MAC proposal today, which one would you pick (from the list above)?

2nd question: Ignoring time-bounded service for the moment, which protocol would you pick from the basis of "moving data around efficiently?"

Tom Baumgartner proposes the 3rd question: "Write down which protocols you find intolerable for any reason - i.e. you couldn't live with adoption of that protocol." More than one is acceptable.

Next thing: start going through the papers. When we break up into groups to deal with the issues log, let's think about which issues are "key?" For instance, which issues are "philosophically at the core" of each proposal, and push for decision on these. I know this is going to cause lots of controversy, heartache, etc. But I think it's better to identify these and work on them, rather than nibbling away at less fundamental issues. Yes, I'm going to press hard on this group to make some difficult decisions.

Comparison against MAC Criteria, P802.11-93/96, by Wim Diepstraten, 93/96

Asynchronous service:

WMAC provides full 802.2 LLC service with unicast/multicast capability.

Broadcast/multicast reliability depend on medium load.

LLC-1-service-only mode without Broadcast available for extreme low power devices.

Isochronous service:

Available as an option, and can mix with stations that do not implement the time-bounded service.

Supports mixed voice/data services.

Higher speed or video support depends on PHY speed.

BSS overlap dependent on PHY channel isolation. When you have a single-channel PHY, you will be limited to a single BSS in a given geographic area.

question: Why is the BSS overlap issue only for the isochronous service?
Wim: In an asynchronous environment, bandwidth is shared "statistically", providing (gracefully) degraded operation for both BSS's. With isochronous service, there is no collision resolution between multiple access points for isochronous service.
Don Johnson: With isochronous voice, cell placement is situated to minimize cell-to-cell interference.
Greg Ennis: This is an issue only in the case of a single-channel PHY. It does not apply to multi-channel PHYs, including hoppers.
Wim says (differently): It will have the same limitations as, for instance, the IBM protocol, given the same constraints.
Jim Schuessler: I think you would admit that you would *try* to channelize, even in the case of asynchronous service.

Supported configurations:

Infrastructure-based network: (async service)

Default is that all traffic via AP, but direction station-to-station also possible.

Single NID per ESA, and unique APID per BSA.

Multiple BSS's can overlap using the same channel.

Infrastructure networks can overlap with ad-hoc networks on the same PHY channel.

Services provided by the AP plus Distribution System:

AP contains "forward function". We call this an "expanded bridge" functionality.

R(e)-Association service provided.

"Initial" and "previous APID" parameters exchanged to support re-association across routers.

Kerry Lynn: When you talk about the expanded bridge function, do you envision bridging with only a "couple of bits" delay?
Wim: No, it's a packet-level store-and-forward function.

Station is responsible for deciding to initiate re-association, based on signal quality.

There are synchronization and power management facilities provided in the access point.

Wireless AP configuration possibilities:

Single AP without Distribution System connection, to expand range.

Allows wireless AP interconnections (wireless infrastructure)

Ad-Hoc Networks:

All traffic is direct station-to-station.

Would use a separate NID and fixed APID (specific to ad-hoc networks)

Can overlap with infrastructure network sharing the same channel.

A station can be associated with an infrastructure network and an ad-hoc network at the same time.

unidentified: "I have a problem with combining ad-hoc and infrastructure AT THE SAME TIME." What does this mean?
Greg Ennis: I think we may be making more of this than it deserves. For instance, if you were hosting a meeting in your facility and people from outside attend, this would allow you to create an ad-hoc network for the purposes of the meeting, while allowing residents of the building to maintain access to the building infrastructure.

Synchronization and power management provisions included.

Stations will listen for a 'timing coordinator' , and will assume this function when none is found - i.e. they will "take on the responsibility and perform" this function.

Infrastructure Network (isochronous service):

Can operate with mixed isochronous/asynchronous devices.

Isochronous traffic in multiple overlapping BSA's require sufficient channel isolation (limited recovery possible).

Reserved but unused Isochronous Bandwidth is available to asynchronous stations.

Need separate "time bounded" distribution system to meet time constraints.

Supports variable bit rates (e.g. ATM backbone, video).

"wireless segment" delay depending on PHY speed <= 25 msec.

Random service sequence provision allows for time in Station to "scan for a better AP" during each frame interval, when "link quality" is deteriorating.

Same provision allows for power management (i.e. scanning for better AP also allows choice of power level, based on path loss to selected AP).

Ron Bjorklund: In this proposal, how do you guarantee that "multi-media" information streams will stay synchronized (i.e. voice with video)?
Wim: The access point controls allocation of isochronous bandwidth, and thereby rejects requests for bandwidth beyond that which is available.
Greg Ennis: There really isn't contention during the time-bounded period of operation, since bandwidth is reserved.
Ron: But there *is* contention in multiple overlapping networks.
Greg: Yes, this is true. And that's the big issue.

Performance Characteristics:

Basic protocol is CSMA/CA with ACK.

Short delay, which is good for bursty traffic.

Efficient and stable throughput at high loads.

Shares medium "automatically", without additional control overhead.

ACK allows MAC level recovery from lost packets.

Overlapping BSA's:

Medium is shared when the same channel is used.

Sharing medium while multiple FH channels collide.

TX power control provision allows improved spectrum re-use.

Mulbiple BSAs cannot overlap the same channel for Time Bounded services.

Performance over Distribution System:

Depend on DS access delay, speed and number of segments.

Performance may depend on power save mode usage.

Slide 13: Idle Power Consumption example

Jim Schuessler: These calculations are for "receive only"?
Wim: This is *idle* power consumption (not necessarily even *receiving*)
Jim: Then I don't know why you even include packet size in your analysis.
Wim: They affect latency and therefore the idle power consumption

Slide 14: Multiple PHY support

Jim Schuessler: Basic synchronization for FHopper already present: Are you suggesting that you assume it exists in the PHY?
Wim: No, the MAC has a timing synchronization, which can be used by the *PHY* for hop synchronization. I don't see any limitation for this protocol in a FH environment. I'm not really sure about the IR environment - haven't thought about it. Feedback welcome on both.
Tom Baumgartner: If you can do it in radio, you can almost certainly do it in IR.

Slide 15: Multiple PHY support (continued)

Note: WMAC being proposed to ETSI RES10 as well.

Bill Stevens: Is the PHY group aware that this MAC proposal plans to tell the PHY when to hop?
Wim: It has been a consensus for some time that it is the responsibility of the MAC to tell the PHY *which* channel to tune and *when*.
Tom Baumgartner: Yes, but there is a vocal minority who continue to object to this.

General

Dave Bagby: I've been asked for the "raw data" collected earlier in the meeting. Before I stick this chart up, I want to caution people against taking it too seriously.
Francois Simon: Did you put your votes in? (Dave's).
Dave: I put my votes in, yes.

Strawpoll displayed as histogram.

Simon Black suggests we drop those proposals which got only "hate" votes.

Dave Bagby suggests that Simon Black inform the authors of those proposals of this decision...

Kerry Lynn suggests we rephrase the intent so that we choose a "short list" of proposals upon which to concentrate further work.

Dave Bagby is loathe to formally drop any proposals at this time. He believes that some of the "less supported" proposals will fade away of their own accord.

There is a suggestion to publish the numbers, which might therefore encourage the "less supported" proposals to be further promoted in written submission and presentation.

Vote: Do you want to see these numbers published in the minutes?

YEA: 20, NAY: 8, ABSTAIN: 8

Vote passes: Poll results will be included in the minutes, however Dave adds that he will include the appropriate disclaimers as to the informality of this poll, and that its results should not be taken too seriously.

(**editor's note:** the minutes as received from the sectretary did not include the results of this poll. However it was reported in the full working group minutes for September 1993, Document IEEE P802.11-93/166 in the AMC subgroup report. Please see that document if your are interested.)

Data Compression in Wireless LAN, P802.11-93/144, by Robert Lutz

Slide 1: Title.

My main purpose in making this presentation is to persuade you to keep data compression in mind while writing a MAC standard.

Slide 2: Transfer Time

(wired vs. wireless)

One of the main benefits of data compression is to reduce the transfer time, by increasing the throughput of the wireless network by use of data compression.

Slide 3: Lossy vs Lossless compression

Lossy: audio/video- can tolerate loss of detail

Lossless: Used for data where loss and/or corruption cannot be tolerated.

Slide 4: Typical Compression Ratio

(graphic showing various compression forms and their performance)

Dave Bagby: Are these compression ratios based on large transfers?
Robert Lutz: Yes, compression ratio is affected by the amount of data over which compression is performed.

In general, the more data you compress, the better the compression ratio.

Wim: This applies only to compression techniques which use memory, right?
Robert: Yes, this applies to *adaptive* algorithms. Other algorithms which start from a pre-built dictionary do not *learn*. Their performance is static, regardless of the amount of data it operates.
Jim Schuessler: You're showing it reaches its asymptote in 2Kbytes (see chart).
Robert: This example is using an LZS algorithm, and is based on a 2K dictionary.
Jim: By the *slope* of your curve, you're implying that 2K is optimum.
Robert: No, this is an example of a *particular* implementation, which uses a 2K dictionary.
Wim: Does the dictionary need to be transferred?
Robert: It depends on the algorithm. For this method (LZS), the dictionary is built on the fly by the receiver.
Don Johnson: In the case of an image, does it matter whether the image is b/w or color?
Robert: We've found that the more detail an image contains, the less compressible it tends to be. This is because the more detail, the more "randomness" the image contains.
Dave Bagby: Is there a connection between the basic structure of the information and "byte" boundaries?
Robert: Yes, there can be quantization effects, for instance if the basic structure of information in other than byte-aligned units, the compression efficiency may be (less?)

Ron Bjorkund: Any compression you do would need to be done *before* any encryption is done.
Robert: Yes, absolutely true.
Wim: You mentioned that this was a curve for a history of 2K bytes. What would be the change in the curve for different histories?
Robert: (draws an example 4K dictionary result), showing that the ratio climbs *after* the 2K point is reached, but not before. The gain goes flat at the 4K point.
Ron Bjorklund: This graph assumes that you are compressing while you are learning?
Robert: Yes - this is how the LZS works.

Slide: Compression Ratio with Small Packets, Resetting History

(graph shown). The compression ratio curve is drawn as if it were a "sawtooth wave", where the compression ratio climbs over the length of a packet, and is reset to zero (i.e. 1:1) at the beginning of the next packet.

Slide: Compression Ratio Without Resetting History

This yields much better results, since history builds over multiple packets. But this requires absolute first in-first out sequencing. Otherwise the receiver can't track the sender in building its dictionary.

unidentified: If you assume you compress above the MAC layer, then compression isn't a MAC layer function (or concern). Have you considered the problem of when a lower layer determines that compression isn't worth it (or is actually detrimental)?
Robert: It could be an issue, or not, depending on the algorithm.

Slide: Multiple Compression Histories

unidentified: Are you suggesting that broadcasts are not compressed?
Robert: No, but just that broadcast packet compression doesn't keep a history (i.e. each packet is compressed as if it were the first, and last, packet ever broadcast)
unidentified: In a multiple compression history environment, this implies lots of memory in access points, and you're also assuming an error-free (or always error corrected) reliable link between the stations (and this is at the *MAC* layer). This is worrisome.
Dave: If I take this approach (i.e. multiple histories), I will need to transfer histories to the "new" access point.
Robert: Not necessarily. You can reset the history anytime you want, including reassociations.
unidentified: As you go through these issues, you become aware of many problems with doing this in the MAC layer. So I ask, why would you want to do this in the MAC?
Robert: Because wireless networks are slower in general than wired networks.
Dave B: I think I have some idea of the problems inherent in the approaches you've talked about. Do you have any "magic tricks" to tell us about?
Robert: Yes. Patience please...it is forthcoming.
Don Johnson: Wouldn't this be good for storage as well?
Robert: Yes. That would be ideal. It doesn't look like it's coming real soon now, though.

Slide: Compression Ratio Due to Dilution

Dilution: As you transfer data from multiple applications, the lower layer compression/decompression algorithm will see different types of data interleaved, causing the history to be "mixed" with different data types. This will cause the compression ratio to fluctuate up and down, in step with the "mix" of data presented to the compressor.

unidentified: Can it be negative?
Robert: Yes, of course. In fact, it always starts out negative (because there is no history).

Slide: Bandwidth vs Traffic

(graph)

The correlation is made that compression relieves pressure on the network by stretching out the point at which the network is "max'ed out" (e.g. CSMA throughput curves).

Slide: Compression Hardware Benefits

reasons why hardware assist is an important component of the overall compression approach.

Q&A at end of presentation.

Wim: You are talking about streaming operation, vs. block operation, right?
Robert: Streaming algorithms are best suited to one-pass compression, almost real-time. Block algorithms allow multiple passes, which can conceivably provide better compression ratio.
Dave Bagby: So far, you've convinced me that I don't want to do compression in the MAC layer. Firstly, compression has to happen before encryption. Since our encryption is going to be done by 802.10, it will be done *above* the MAC layer. Second, many of the efficiencies are based on multiple-packet histories, and this is very difficult to achieve (if it can be done at all) at the MAC layer.
unidentified: I may argue the security argument. I think we previously argued that we would combine encryption and compression as a single function.
Dave B: Yes, but still this is done above the MAC layer (i.e. 802.10).
Leon: The way we've shown security is logically above the MAC, but where it actually gets done isn't necessarily constrained in this manner. Perhaps the layer management allows us to tell the MAC when and how to do these things (encryption/compression). (this sounds like delegating the function to the MAC, even though it isn't within the MAC's domain).
unidentified: In deference to Dave's thing about .10, I think you'd find that anything registered in .10 as an algorithm, you could also register methods which included encryption.
unidentified: What compression ratios could you expect if you assume you reset after each packet (assume, for instance, a 500-byte packet length)?
Robert: I can't predict, because it is packet-length and data dependent. You could run representative data, though to get a feel. With 576 byte packets, I would guess 1.5:1 ratio.

Tomorrow morning, we will restart with the next compression item.

Meeting adjourned: 4:57 PM

Tuesday, September 21, 1993

Meeting called to order at 8:43 AM, by chairman Dave Babgy. Bill Stevens secretary (edited later by Carolyn Heide).

Analysis of a Data Compression Method as a MAC Option, P802.11-93/158, by Waychi Doo

LZ77-based compression program tested using four types of files. Compressed into blocks; history buffer flushed after each block.

Several slides presented, showing compression ratios obtained for each file type:

- ASCII text
- binary file

- bitmap image
- Image (TIFF) file

Compression ratio drops when: data is very random, non-correlated frames (images?), block length is short.

Slide presenting a number of conclusions:

Small frames don't compress well - larger MSDU size desirable and/or segmentation and reassembly support

Non-correlative frames

Don't compress control messages (they're too small)

Separate history buffer for each MAC destination

When to flush history buffer? (each frame or several frames)

Negotiate data compression during association

Convey indication of compression status in MAC header

Q&A:

Dave Bagby: You mention segmentation and reassembly. Is it worth doing if you don't segment and reassemble?
Waychi: Segmentation and reassembly adds a great deal to the performance of compression.
Dave B: Segmentation and reassembly implies a reliable, ordered delivery, which (if done at the MAC layer) is beyond the traditional scope of a MAC service.
unidentified: It seems that data compression should be done at a higher layer.

Frame Prioritization in a CSMA/CA MAC protocol, P802.11-93/159, by Rick White

This is based on the NCR approach of using the inter-frame gap. They have certain priority levels, and the length of delay after "busy media" is related to the priority of the frame to be transmitted. In other words, higher priority frames will be transmitted "earlier" (thus capturing the medium "earlier", causing lower priority frames to defer).

Wim and Rick discuss the nature of the "backoff algorithm", ensuring they agree on what this means. Others join the issue. A future paper from Rick will illuminate this topic.

Slide presenting a possible mapping of frame types to priorities (e.g. data, acknowledgement, point coordination function, asynchronous, physical layer)

Slide presenting an example of three data frames, followed by an ACK in the reverse direction, illustrating different priorities of packets.

(general questions about the example.)
Leon Scaldeferi: Wouldn't you want the ACK to be at a higher priority, so the feedback was as short as possible?
Rick: Well, the gap between the three data packets is *very* short, since they are back-to-back. The ACK is the highest priority other than this.
unidentified: It seems that what we want is for the delay before the first frame at a *lower* priority, then subsequent frames (in the burst, so to speak) would occur with shortest delay. This looks like a form of segmentation and reassembly.
Rick: Yes, in fact this is what I mean to be showing here.

Don Johnson: Do you go into the backoff algorithm, then, going back into the asynchronous contention?

Rick: Yes, that's how it works.

(Note: Rick confirms that the example of three data packets plus ACK *is* in fact, an example of segmentation and reassembly)

Point coordination function is set at priority 3, which includes time sync beacons to synchronize frequency hopping systems.

Wim: Is it required that all stations receive this beacon?
Rick: Not really. Each station can keep a count of how many periods have elapsed between beacons, and can therefore miss several beacons and still assume it is in sync.
Rick: For instance, say a data frame is being transmitted at the point in time when a beacon should be sent, the beacon must be deferred, so...

Slide 8: Dwell Stretching using Prioritization.

unidentified: Would you then shorten the following hop period correspondingly?
Rick: Yes; this also helps meet FCC dwell requirements
Jim Schuessler: Doesn't this introduce significant "jitter" in delivery times of time-bounded services?
Rick: You can control this by setting upper bounds for "stretching" and enforcing it through the use of segmentation.

Slide 9: Time Bounded Services Requiring Different Periods

You may have time bounded services of different periods. Slide depicts two time bounded services at different periods (e.g. 3ms and 5ms)

Wim: You didn't indicate any time bounded priorities in your priority scheme. I see a PCF, but I see it being used to transmit a beacon signal. What priority would you use for time-bounded packets?
Rick: I would use this same priority (PCF) for time bounded packets.
Tom T: Would access points obey the same priority as mobile stations?
Rick: Yes.
Jim Schuessler: Given the priorities you have for a given transmitter, isn't it possible that, given large segmented frames, (subsequent frames with "high" priority), wouldn't this push out time-bounded services quite a bit?
Rick: There would be some jitter, yes. I think it can be adjusted with parameters.
Jim: Perhaps we need to give time bounded services a higher priority.
(a discussion about what practical period times may be for time-bounded services)
Leon S: ATM uses a 6ms time period for voice. That is considered acceptable.

undentified: What kind of dwell time are you assuming?
Rick: About 20-30 ms.
Wim: (observes) This looks the same as our (NCR) proposal, with the addition that you can support multiple segments of a (long) frame in a single transmission, with a single ACK.

Slide 10: Transparent Relay

Wouldn't Station "a" have a variable delay in receiving an ACK (based on whether or not a transparent relay is occurring)?
(lots of discussion -unstructured)
(the slide shows an "asymmetric" packet flow, where B can't hear A directly, but A sends an ACK which is *assumed* to be receivable by A - this leads to lots of confusion, which discussion attempts to resolve)

Slide 11: Priority 5 - Physical Layer

Packets sent at the physical layer are "transparent" to the MAC (i.e. the MAC never sees this traffic)

Tom T.: Would the MAC transmit these frames?
Rick: No, the MAC wouldn't know about it.
Wim: This would contend with the asynchronous frame contention, wouldn't it?

Slide 12: Conclusions

Wayne Moyers: Could you describe the "initial" state of a starting-up network.
Rick: If it's at startup, nobody is sending, nor waiting to send, so the access point will send the first beacon without contention.
Wayne: So you gotta walk through and listen on all of the various hop groups until you hear a beacon?
Rick: there are a number of ways to achieve synchronization
(random discussion)
Simon: What is your vision of what you would do with the PHY layer messaging?
Rick: To relieve the MAC of some detail, hopefully easing the task of making a PHY-independent MAC
Kerry: Try substituting the word "Management" for "Physical", and I think you'd have an easier time of it.

The Necessary Conditions For Support Of Connection-Type And Connectionless Services On Wireless LAN, P802.11-93/xxx, by Chandos Rypinski

Presumption. All information that is discerned by a receiver must be in the "bitstream". Other forms of information, such as "carrier present", or "signal strength" can't be depended upon.

Ability to render connection-type service on an asynchronous packet medium is dependent on reservation of future transmission space/time in the medium to assure timely delivery of each octet bundle.

- There is no reservation algorithm unless the results of it can be used to control and schedule use of the transmission medium. The resulting requirement on the PHY and MAC is the ability to transfer.
    1) the information on which access decisions are made, and
    2) the enabling instructions for station originated transmissions.
- The scheduling algorithm must achieve some performance parameters infrequently used in specifying packet LAN:
    3) worst-case access delay, and
    4) transfer delay which is the sum of access, quantizing and propagation delays.

Delay:

Delay is limited by echo from reflective 2/4 wire junctions or acoustic coupling in 4-wire handsets. Tolerable added delay from a local distribution system is an argument on how to share about 30-40ms of tolerable delay with other parts of the public network. The important sources of delay are:

a) quantizing delay -- e.g.; 6ms are required to accumulate 48 octets of samples for a 64kbps channel. this example corresponds to B-ISDN (ATM) voice transmission.

b) coding delay -- codecs for speech compression generally require 10-60ms between input and output. Avoidance of this delay is inducement for not using highly compressed voice.

c) access delay -- the worst case delay before transmission on the medium can start after the assembled octets are ready for transfer.

d) transfer delay -- the sum of the quantizing, coding (if used) and access delays and the added delay between the start of transfer at the originating end and the start of output of transferred octets at the receiving end. (on cross-continent and overseas connections, the propagation delay is significant.)

In general, the sum of these delays and those of the tandem telecom facilities must be less than 30-40 ms.

Don Johnson: What delay could the user tolerate in the absence of echo?
unidentified: 400ms, according to the CCITT.
Chan: That's true, from an echo-only perspective, but delays of this length show up in the "turnaround" when one person wants to talk too early (not accounting for the delay).

Q&A:

Ron Bjorklund: Would you view the "connections" between each end-user station and the access point as a point-to-point connection, with no contention with other stations?
Chan: I view the associations between end-user and access point as a multidrop network, where the access point serves as a "concentrator", with typically 10-20 stations per access point.
Tom T.: (initiates an inquiry which results in a "centralized vs distributed" issue)
unidentified: The functions must be performed. Where they are done is an implementation decision. But the functions must be performed nonetheless.
Kerry Lynn: Should we consider it's valuable or necessary to support multimedia over 802.11, I would suggest that there is a healthy market for networks which do *not* support multimedia. Earlier, you said that most proposals on the table would not meet your criteria...leading to the assumption that they should not be considered...
Chan: While I say they do not meet these criteria, I wouldn't say that those of a different persuasion shouldn't be allowed to pursue their goals; Likewise, those "other" persuasions should see the case for (this) approach. There is an implicit assumption that "higher speed" means "higher cost". This equation does not necessarily hold. High volumes could reduce high rate system cost as well as reducing the "rate" of the system. Also, "scaling down" is easier than "scaling up". Therefore the higher rate system may more easily be scaled down to meet less aggressive goals than the converse.
Kerry: Assuming, of course, your infrastructure approach...
Chan: I don't see any hope of ad hoc groups operate in the presence of an infrastructure (in an organized fashion). The beacons generated by the infrastructure could (should??) be used by ad hoc stations to synchronize their operation with the infrastructure.

Comparison and Commonalities of Asynchronous Sequential Access and Adaptively Partitioned Periodic Frame MAC Proposals, P802.11-93/163, by Chandos Rypinski

OVERVIEW OF COMPARED PLANS

Two series of contributions on MAC have been presented to 802.11--ASA and APF.

The two plans are quite similar in objectives and asserted functionalities but differ considerably in certain aspects of implementation. With modifications, either plan could meet the same end functional objectives.

Both plans have in common:

contention on request

exclusivity of channel use for all subsequent transfers

division of packets into shorter segments for transmission, and

provisions for retranmission of failed segment transfers

While detail changes and improvements have been made in both plans since the cited references, the principles involved appear to have remained constant.

It will be asserted that "asynchronous sequential access" is a better choice.

Slide (very dense) comparing APF and ASA

Chan observes that the IBM proposal *assumes* the use of frequency hopping. He doesn't agree with this approach, but will keep it out of his comparison. He therefore assumes that both approaches are provided with equivalent "media resources", for the sake of comparison.

CONCLUSIONS

The asynchronous access method provides significant advantage relative adapatively-partitioned regular periodic frame structures because of the previously given reasons which are condensed and summarized as follows:

1) There is no fragmentation of unused space as a result of transfer organized use of channel time.
2) The minimum and worst case access and transfer delays for a given set of transmission parameters are smaller.
3) There is minimal built-in assumptions about traffic distribution and characteristics.
4) The definition of the MAC is simpler and easier by an order of magnitude than for any channelized or slotted system
5) The risk of undefined hangup modes is far lower.
6) The ultimate capacity and breadth of function of a given frequency space is greater.
7) The single channel wider band radio will be simpler, easier to make frequency independent.

A few years ago, I got my first big hard disk, and said, "Hallelujah! I can now segment this disk into several partitions, and store my data according to category. After several months, I discovered that my free space was *also* partitioned, such that the free space in any one partition was insufficient for the task at hand, whereas if all the free space were contiguous, it would have been sufficient."

The same holds true for the wireless LAN. I admonish all of you against slicing up the radio resource into fixed partitions, thus perpetuating this disadvantage.

Don Johnson: It appears that the *whole system* must be coordinated. This implies that it will be difficult or impossible for any users which are *not* a part of this "system" to operate.
Chan: I assume in my system that burst transmission is used, and coordinated by a central control point. If you take your voice packets and simply transmit them as soon as you can, and the bandwidth reservation system (described previously) does its job, there *will* be sufficient capacity for all traffic. Now, if you have unrelated groups (nearby), it does not follow that it "won't work"; the unrelated groups will result in a reduction of capacity (i.e. it will just be "interference")
unidentified: I've heard a wide range of "ranges" being discussed; You speak of "a floor in a building", others speak of areas that you could "drive through." This is a very wide variation.
Chan: I don't think you've heard that much in *this* venue. If you're traveling in the PCS community, you certainly will hear about vehicular ranges. But not in 802.11.

(Chan tells a story about consulting a cellular operator on "cell splitting". His expert consultation resulted in his getting fired. The subsequent consultant(s) eventually convinced the operator that (Chan's) proposal to multiply the infrastructure investment must be taken seriously...)
Kerry Lynn: I think maybe some of the work we put into designing a MAC should be put into figuring out how asynchronous(only) system could be made to peacefully coexist with time-bounded systems, when they come to exist. Makes the observation that an infrastructure system requires an up-front investment (in the infrastructure) which becomes cost-effective only in the "end", when the user population grows to a sufficient size.
Chan: Communication has a value to the organization and the users. I'm talking simply about value. So if I said the more access points, the more valuable to the organization, perhaps this is a more accurate description of the economics. My proposal includes an autonomous ad-hoc mode, which is a subset of the whole system design. So, though it appears at first glance to be "ponderous", in fact my proposal "grows up" from ad-hoc to full infrastructure, over time. So nothing has to be thrown away. If we look at the economics from the view of "avoidance of early obsolescence," I think the merits of my argument increase.

I realize that not many people are in agreement with my views on high capacity emphasis, but I think it's a big mistake to not take these matters into account.

Kerry Lynn: You may want to note that there are people who agree with your view, but don't agree that this is the band in which to do it...
Chan: thank you.

In today's economics, "logic" doesn't cost very much at all. So it does not follow that (logic) complexity results in higher cost.

On the subject of "avoiding hang-up modes"... I believe that this system is inherently more immune to getting into "undefined states" than many other proposals. As a rule, the longer a "state" exists, the more likely the "real world" will change underneath it, so that it no longer agrees with the logic's assumption of the current "state." Minimizing the number of states in the design is an important objective.

Being inherently a "single channel" approach, this system design matches equally well with any radio bandwidth and/or frequency band. Such changes are totally unrelated to the rest of the system design.

Also, (though possibly counterintuitive) it is possible for a wideband radio to actually be cheaper than a narrowband radio. This is so because much of the cost of a radio is related to its *precision*. So, do not leap to the conclusion that "wider" is "more expensive".

Revisit To Channelized MAC Proposals, P802.11-93/1xx, by Chandos Rypinski

Reviews previous contributions... All of his proposals are common in the use of message-based handshakes.

Discusses a tiling of 9 squares, where you would use 9 frequencies - one per "square". Then you would need a common "hailing frequency" which would be used to coordinate users with one of the 9 frequencies/squares.

The channelization had no visibility to the MAC (i.e. the MAC was not involved beyond maintaining a channel number, which was given to it by some other layer or function (management?))

If you believe the process of listening for a clear channel is an unreliable process, it's better to just "throw it out there" and use an immediate acknowledgment to confirm that a path exists.

The most telling argument for the complexity of channelized vs single channel, just measure the thickness of the paper proposals for both; the channelized system is usually twice as thick!

General

Time to go into small groups. We will create four small groups. Here are the suggested topics:

Group 1: Leon Scaldeferi recording
- 6 - Security (6.10)
  wire equivalence algorithm brainstorming
- 9 - Compression
  9.5 Yes/no?

Group 2: Jeanine Valadez recording
- 4 - Network topics (ad-hoc and infrastructure)
- 4.5 - simultaneous ad hoc and infrastructure?
- 10 - Coordination & 11 - access points

Group 3: Jim Schuessler recording
- 15 - services (TB)
  -15.8 all support Time Bounded?
  -to simplify task, could we leave TB to later given hooks in proposals on the table?
- 18 - Data Rates
  - 18.5 is data rate agility only a phy matter?
  - MAC required to handle multiple simultaneous PHY data rates?
  - explore each MAC proposal re this ability.

Group 4: Arnulf recording
- 17 - addressing (extent, support)
  - explore issues, flesh out arguments.
  - see doc 93/1136 for some input

Group breaks into subgroups for the rest of the day.

Meeting adjourned: time unrecorded.

Wednesday, September 22, 1993

Meeting called to order at 8:45 AM, by chairman Dave Babgy. Bill Stevens secretary (later edited by Carolyn Heide).

First order of business is to get the reports from the breakout groups from yesterday afternoon - see if it generates discussion in the larger group, and see if we can generate any official actions as a result.

Group1 - Security and Compression

The group dealt with two issues - 9.5 and 6.10. First thing is that the "compression", as written in the issues log, really needs to be reworded. (Francois will take care of this)

What we did was look at a couple of techniques that were described the day before, identify some of the characteristics you may be interested in.

(table projected, showing matrix of three types of compression, with four characteristics per type)

Then we looked at security (Table showing privacy algorithm and compression algorithm)

Issue 9.5 "Shall 802.11 provide data compression at the MAC layer?"

Recommendation: "yes". (Even if modest compression ratios of 20% would be valuable.)

Raised two other issues:

Wim: Did you factor in the increase in complexity of adding compression?
Leon: If you look back at the previous table, you see we judged the complexity of each of the several approaches for compression. Some of the approaches had low impact; others had higher impact.
Wim: But isn't this a relative judgment of the complexity of one compression approach vs others?
Leon: It was also meant to represent the impact on the *system* complexity as well.
Wim: What about impact on delay, due to compression processing time? For instance, block compression schemes might require you to preprocess a packet, resulting in an up-front delay. Also, the need for a length field at the front of a packet, you will need to "pre-compress" before transmission is begun (thus, an up-front delay). Robert Lutz (Stac) agreed that if you have a length field whose value depends upon the result of compression, this is certainly true.)

Issue 9.5A How should compression be supported and specified?

Alternatives: Yes/no field for compression and a unique algorithm identification field.

Issue 9.5B Should the default also be "none"?

Recommendation: "Yes"

Dave B: Comments (personal opinion). Something doesn't jive quite right here. The desire here is to get a more efficient pipe. (especially since our pipe is narrower than most). But even with what you've answered here, I can't convince myself that this is a MAC issue. It makes sense to say "it would be nice to have this capability", but I have a hard time conceptualizing that the MAC would decide to turn this on and off, and I also think it lives conceptually at the same level as 802.10 security. Perhaps it is better to ask the exec committee to create a new PAR (and therefore a new group) to address compression.
Leon: there are some management functions that have to come down from some higher layer to *control* the activation of these features. There is a feeling in the group that there the increases to be gained by compression that are significantly valuable.
Dave: What hangs me up in this is the "in the MAC" phrase. From the more formal standards viewpoint, I just go "tilt" when I hear "in the MAC." 802.10 used to be a thin line between LLC and MAC, and now we're making this line thicker (by adding compression?).
unidentified: Whether or not it is done in the MAC, there is a case to be made for *control* of it in the MAC, (i.e. switching it on a packet-by-packet basis).
unidentified: What mechanism does 802.10 provide for identifying which packets are encrypted vs not? Could the same mechanism be used to indicate the presence/absence of compression?
Ron Bjorklund: It was felt by the group that, even though these features (compression, security) are valuable to *all* forms of 802 LANs, it is felt that these features are *most* valuable (and important) for wireless LANs.
Kerry : Since we're including 802.10, why not include compression, since everybody would desire "higher bandwidth?"
Joe Kubler: For those compression types that can "stream," it makes better sense to do compression in the MAC. Otherwise, you almost guarantee a serial delay (pause) while waiting for compression to be done. (someone comments that this is an implementation issue)
Wim: Could you "bind" compression in with encryption, such that specifying an encryption "type" (in 802.10) would also specify a compression type?
Leon: I think that could work.

Dave B: We have decided that we are using 802.10 to do security, mainly because it *exists*, and we really can't go off and do duplicate work. It seems to me that compression needs to be done *ahead of* encryption, if you make a "chip" that does 802.11, you have prevented the inclusion of encryption (unless, of course you put it in the 802.11 chip also). It seems that it isn't possible (reasonable?) to do this as an 802.11 thing.

Dave: I do like the approach of keeping them (encryption/compression) separate. I may choose different encryption methods, based on the level of security, and the level of compression.
Bill Stevens: It seems to me that those uses of a LAN which most benefit from compression are those where streams or large blocks of data are being moved *en masse*. (e.g. file sharing, file transfer). I observe that it is actually rare for application programs to perform these actions at a low level themselves - usually they are performed by network operating system services. Just because this committee wants to see compression *done*, doesn't automatically imply that it should be done by *this committee*. Unfortunately, there is no equivalent "network OS" standards committee with which to liaise on this issue.
Francois: We did decide on some functions that can be enabled/controlled in the management "sidebox".
Dave: There seem to be two choices : a) Yes, we should do these things ourselves, b) Yes, but we should find some other place to get these things done. I don't hear anyone saying, No, we shouldn't be doing these things at all.
Leon: The intent of answering "yes" to Issue 9.5 was to force the exposure of 9.5A and 9.5B, which would not be explored further if 9.5 were to be answered "no."
Dave: We should try to get 802.10 to adopt the issue of compression. If they won't accept the charter, we must revisit it. But to me, it still makes better sense for it to be done there (above the MAC), than within 802.11.
Leon: The model we currently follow shows .10 immediately above the .11 MAC, but this doesn't imply that they must be physically implemented that way. Perhaps
unidentified: If the layering defined by 802 is wrong (for us), we should be lobbying to change it.
Robert Lutz: Isn't this just semantics? (YES!) Why can't you "suck" 802.10 into our MAC?
unidentified: If it (the layering) is broken, let's fix it!
Dave: I'm saying it ain't broke.

Dave: I'm going to take you from architecture to implementation for a moment. If you try to do the 802.11 part in hardware, and then you try to do the security part in software, I think you're hosed.
Kerry: One of my tests of a good architecture is that it can be applied in ways its creator didn't anticipate. If you try to combine N times M combinations of encryption and security, yes, you will eventually break the model. But I think it isn't unreasonable to go to 802.10 and ask for a small set of "numbers" with which to combine security and compression types, and this will take care of it for us.
Dave: The only problem I see is that you lock yourself into only a small number of possible options - it isn't open-ended.
Leon: I think you need a separate flag to allow selection of a compression algorithm (to separate it from the encryption).

Motion #1: **Compression will be supported by providing the necessary mgt hooks in 802.11 to invoke the mechanism of 802.10 and we will formally request that 802.10 extend their work to include support for compression. If 802.10 declines, we will have to revisit this topic.**

Moved by: Leon
Seconded by: Tom Baumgartner

Motion Discussion:
unidentified: Assuming we vote YEA on this, when will .10 be informed, and how soon can we expect to hear back?
Dave: I will inform Vic within the week, we should expect .10 to hear from us formally within the meeting in November.
Leon: If Vic talked to the chair of .10 following this meeting, I would expect to hear from them within the Nov. meeting.

Approved: 24 Opposed: 3 Abstain: 5 ***Motion #1 passes***

Report from Group 2 - Coordination Function Topics

Issue 4.5 "Can a station be a member of an ad hoc and non-ad hoc network at the same time?"

We concluded:

- Routing and bridging is *not* the issue.
- People will do it anyway (using sleep mode)
- Member would have to have two net IDs
- Conclusion: Standard should be mute on this issue; the standard shouldn't require simultaneity; not doing anything to explicitly enable or disable this function.

New issue: dealing with multi-tenant configurations

Motion #2: **To adopt these recommendations**

Moved by: Jeanine Valadez
Seconded by: Dave Voth

Motion Discussion:
Charlie: Do you have minutes that show more detail on how these conclusions were arrived?
Dave B: Notes were taken, which are given to Francois for inclusion in the issues log. But no, official "minutes" were not taken.

Approved: 26 Opposed: 0 Abstain: 3 ***Motion #2 passes***

Issue 10: What Coordination Function will be specified?

Propose change to CF definition: delete "and receive"

Major finding: distinguished between PCF and POINT/CENTRALIZED-CONTROL FUNCTIONS; e.g., pwr mgmt, store & forward, DS access, channel opt, network planning, these have nothing to do with CF

Concentrating focus allows better comparison between PCF and DCF

Reviewed pro/con matrix

Call to action to corroborate claims

- 2.3, 2.5 should be deleted because N/A
- also propose delete 1.4, 2.8
- add: DCF Pros: better to deal with other Tx's in BSA; well suited to async traffic

- add: PCF Con: doesn't work with single-channel PHY in overlapping BSA's
- change: 2.6: "is higher" to "can be higher"
- change: 2.7: "is required" to "is better suited"

Recommendation: PCF or DCF? Both (a "soft" recommendation)

New Issue: Store and Forward provisions in ad-hoc net?

Francois: This should be recorded in the issues log as "added arguments", not "change" the wording of the issue. This would violate the "history" of the issue.
Phil: I hope we can preserve the point of how we desire to change, so this input isn't "lost" by editing it into conformance with the structure of the issues log.
Dave B: I didn't mean to be so strict in how it gets recorded; I will trust the issues log caretaker (Francois) to get it recorded faithfully.

Dave B: I saw the group wrestling with the fundamental nature of this issue (i.e. religious in nature). I suggested that rather than try to make a polar choice, perhaps they could list pros and cons for both approaches. Now I see a recommendation to close the issue by adopting "both." I don't think this does any more than "mechanically" close the issue. This draws us to the subject of "how do we narrow things down?" This is something we need to make progress on during this meeting.
Kerry: Are you saying that you're more comfortable with making a call one way or the other, vs "both"?
Dave: I think they did the best they could, given the task, but ultimately we need to work through the dichotomy, and arrive at a better (than either) approach.
Phil B: I think we were pretty clear about the pros and cons of both. The reason we couldn't be more definitive in our conclusion is that we didn't have the conclusions from the group considering "is time-bounded services a requirement?" Also we recognized the bias (for DCF) in this smaller group.
Simon: The strawpoll earlier this week showed support for approaches favoring both DCF and PCF. I think what we should do is use the results of that strawpoll to reduce the number of MAC proposals during this meeting.

Simon: I would ask for a slot on Thursday morning to work on this.
unidentified: Many of the pro/con evaluations were not on PCF/DCF, more on a specific MAC protocol which implied PCF or DCF.
Jeanine: These "functions" which were better supported by PCF/DCF mudded the process of considering PCF/DCF on their own merits.
Tom B: I feel that this issue shouldn't be closed until we do some narrowing the field of MAC proposals.
Kerry: In at least one submission, it is suggested how to migrate from a DCF to a PCF... Perhaps we should open a new issue exploring the existence of *both*.

Motion #3: **To close this issue with the specification of "both".**

Moved by: Jeanine Valadez
Seconded by: Chan Rypinski

Motion Discussion:
Kerry: I would like to amend this motion to further investigate migration of DCF to PCF as part of the "both".
Don Johnson: The progress on this issue would be better to illustrate which functions are better supported by DCF or PCF.

Jeanine: the group investigated the "better suited to ?CF", but due to the group's (DCF) bias made us uneasy on making a "choice."
Dave Voth: It seems that the only way to proceed is to choose ONE of the two, and then make improvements to it (to resolve this matter).
Dave Roberts: I didn't get the impression (in the smaller group) that we expected to come to the "motion" of adopting "both". I think this approach skirts around some of the things in peoples' heads (which this decision doesn't elaborate).

Dave B: I wish to table this motion, until later, so we can finish this report.

Motion #3 tabled

New Issue: Store and Forward provisions in ad-hoc net?

IETF Mobile Networking, P802.11-93/1xx, by Charles Perkins

The IETF is the Internet Engineering Task Force. This has been in existence for ten years or more. Recently the mobile networking group was chartered to investigate the expansion of internet protocols to operation in a mobile environment.

I'll talk about the problem, how it affects internet protocols, and the model of how we expect to incorporate mobility into internet protocols.

Slide showing the many positive aspects of mobile networking they wish to achieve.

Point: We look at the "problem" from a Layer 3 perspective. (as opposed to 802.11, which doesn't include Layer 3 in its domain)

Hard requirements:

- Continuous access (multiple networks)
- Backward compatibility
- Weak security (?)

We're talking about mobile networking. Mobile and "wireless" are not completely the same thing. For instance, if you unplug your computer from one subnet and carry it to another location and plug it into another subnet, this should work as well as if you did the connection wirelessly.

Kerry: If we provide security at the MAC level, would that benefit accrue for your systems as well?
Charles: Yes, and let me note that I am here as liaison from IETF, to make some "weak" recommendations - things that you could do that would help us out. But remember that our approach must work equally on wired (see above example) networks as well (so implication seems to be that they can't depend on 802.11 approaches).

The group is focused on providing mobility while maintaining the same IP address. This clearly contrasts with the traditional thinking, where an IP address IMPLIED a fixed location.

We want to make sure that existing routers and host computers do not require changes to work with mobile computers.

Phil Belanger: Is your compatibility requirement so strong that you would require that a fixed (non-mobile) computer will not even require a software upgrade?
Charles: Yes, in essence that is what we are saying - especially with respect to routers. Now it may be true in some cases that this is impossible to meet (discussion of short timeouts becoming too short - but this is not prevalent. TCP timeouts are long enough to not present much of a problem here).

It is observed that this is a routing problem - the rest of the protocol set is essentially unaffected (e.g. TCP is a reliable stream protocol, which will continue to function in a mobile environment, provided that datagram routing "works").

"Mobile" IP addresses are thought to have a "home subnet", of which it is a member. When a packet is (naïevly) sent out, the "foreign agent" will handle the forwarding for the mobile node. The address used by the sender is known as the "care of" address.

(shows a diagram of "triangle routing", which depicts the inbound and outbound datagram paths into and out of a mobile host)

Security: Today's Internet is vulnerable to "spoofing" (given a router somehere in the path that handles datagrams in both directions between two endpoints).

With our layer 3 solution, we're not guarding against someone in the same "cell" (wireless or physical bus/ring) from "spoofing".

Previous proposals:

- IBM Loose Source Routing
- Sony
- Carlberg's Host Route
- Coumbia (JI) MSSs (mobile system servers)
- Matsushita
- IBM Readdressing
- CMU MHRP (mobile host routing protocol)
- Myles/Perkins MIP (Mobile IP)
- SMIP (CDPD-like)

With CDPD, they 're selling unused "cellular capacity", consequently they will charge for it. They haven't solved the "triangle routing" yet, but are looking to IETF for this solution. The reason they haven't made progress is they haven't figured out how to solve the "billing problem." If you eliminate the "third apex" of the triangle, then they have no way to bill for some of the traffic...

Current IETF direction

- Encapsulation (not options)
  Advantages:
  - Options slow down routers. Encapsulation does not suffer from this option.
  - There are implementations of TCP/IP where the host will crash if it receives a datagram with "unknown" options.
- Solicitation
  Like an "advertisement" (see below), but the mobile host "asks" as opposed to (unsolicited) broadcasting.
- Advertisement
  A "beacon", for instance, is a form of advertisement. One of the things we want the foreign agent to advertise is its "care of" address.
- MH <-> HA Registration
  Mobile host could register directly with the HA (home agent), or it could register with the FA (foreign agent), who could participate in propagating the registration with the HA.

- FA <-> HA Registration
  see MH <-> HA Registration (above)
- ICMP "Remote Redirect"
  The "remote redirect" using ICMP is more reliable than "encapsulated" datagrams, because it is observed that many hosts which may "crash" on a new IP type, will properly ignore an ICMP message of new (unknown) type. So it's "safer" to use this approach.
- Validation procedures
- No help for "ignorant" hosts
  If you have a correspondent host that doesn't understand ICMP remote redirect, then it still works, but with sub-optimal routing.

LAYER 2 INTERFACE

- Cell association events

  It would be good to know about when mobile hosts associate and unassociate with an "access point" (or whatever).
- Carrier detect
- Base station MAC address

Francois: Is the definition of Layer 2 a pure definition, which includes Layer DLC as well as MAC?
Charles: Yes... Perhaps I should be presenting to 802.2 as well. I should make it clear that I am here as a liaison from IETF, and I think I'm being complete and "fair" in my presentation. But remember that all of this is under debate, and so to some degree, what I present is my view.
John: We would like some mechanism to alert us to the possibility that a mobile host may be served from two "cells", and (presume?) have some say in which "cell" provides service (is this what he said?)

SUMMARY

- Mobility solved at Layer 3
- TCP/IP solution can adapt to other protocols
- OS and machine independent
- Framework for mobility is designed
- Applications work without change
- A working group proposal may emerge this year

<< A draft proposal likely within a month >>

Persons interested in more information may :

Join the mailing list by sending email to: mobile-ip-request@parc.xerox.com (read by Steve Deering, so say something "human recognizable" in the body of the email)

Discussion traffic will be sent with a "from" address of mobile-ip@parc.xerox.com

To contribute, send your email to this same address (mobile-ip@parc.xerox.com)

Discussion:
Greg Ennis: There is a particular configuration I didn't see discussed. Are you assuming that the home agent and foreign agent are necessarily physically attached to the "wireless" network?

Charles: I'm not assuming that - there's a lot of latitude here. If you are saying that the FA is itslef a mobile device, then, in some configurations, yes that could work (e.g. in an airplane at 30,000')
Greg: In 802.11, we have a "distribution system" which handles mobility "at a MAC address" level. One configuration we may see is these "access points" interconnected via other 802 wired LANs, so I think we will handle much of the mobility you are addressing here. But we *don't* handle IP internetworking, so your approach complements ours.
Charles: The configuration you outline has certainly received consideration. We want to come out with a workable draft ASAP. If it is the case that our first draft doesn't deal with your case, then perhaps an "outcry" will drive us to work further on the problem. But our emphasis is to get a draft out soon. It can evolve as needed.

Dave Bagby: Adjourns this meeting (due to the late hour), but will allow the discussion to continue "informally".

Meeting adjourned: 12:50 PM

Thursday, September 23, 1993

Meeting called to order at 8:37 AM, by chairman Dave Babgy. Bill Stevens secretary (later edited by Carolyn Heide).

Report from Group 4

Issue 17.2

New alternative: If multichannel system negative ack could use a spare channel for error correction.

Issue 17.5 What is meant by addressing?

Recommend: Use IEEE 48-bit addressing scheme

Discussion:
Chan: It seems unarguable that you have to support the IEEE addressing scheme. But this may not rule out the use of "short" addressing (to save time on the channel). For instance, ACK messages can't be *unaddressed*, but they could suffice with short addresses. I would prefer to see wording that would allow for this. Also, there is E.164 addressing (CCITT), and 802.6 uses this, along with 802 (48-bit) addressing. To me, I would probably be an advocate (in setup messages) of using the 802.6 format, in which case all the definitions for alternative addressing are already in exist (in 802.6).
Leon: You're talking about having to transmit these long addresses on each packet. The use of a short address after "association" (or other "startup" message exchange) can be helpful in keeping down message lengths. A directory of long->short address mapping would need to be kept.
Ron Bjorklund: Is the assumption that all packets use this addressing? If so, this would be costly.

Dave Bagby takes a very informal stawpoll. How many think this issue needs further consideration. Almost all hands went up. It will be left open.

Issue 17.6 Global addresing and Directory service effect on MAC?

The group basically didn't understand what this issue is about. Dave B. would like to discard this issue, rather than thrash about trying to *invent* a meaning for this issue.

Leon: I think this is very close (or identical) to the addressing matters Charlie Perkins was discussing yesterday, and this is Layer 3 addressing--not our domain.

Dave Bagby takes a vote: Shall we close this issue without further consideration?

YEA: 28 NAY 1 ABS 3

This issue is now closed.

Issue 17.7 Does MAC supply packet number to PHY?

Recommendation: MAC does not supply to phy layer packet numbers

Bob: What we are dealing with is MPDU sequence number, etc., all of these are covered by CRC, and so these are MAC numbers. They are not the "business" of the PHY.
Chan: I think we should close this issue for the same reason as the previous issue (above) i.e. it is not within the domain of the PHY.
Jim Schuessler: I would argue to close it for the same reason as Chan.
Ron Bjorklund: My gut feel is that there may be some implications in the area of multiple-packet compression, or something like that. I don't want to see this issue be "blown away" due to lack of understanding as to what it is about.
Rob Lutz: I'm uncomfortable with closing this issue because 1) we don't know what it is about, and 2) there doesn't seem to be anyone present to comment on this.
Wim: What is meant by packet number?
Francios Simon: This was a very early issue. We were somewhat shooting from the hip when we started the issues list.
Simon Black: The way we primed the pump in creating the issues list was a couple of hours of freewheeling brainstorming. I see no need to keep the issue alive, especially since nobody seems to understand its reason for existence.
Ken Ju: We weren't sure what a PHY packet number is. What we thought made sense that the PHY might want something to do with "quality of service", but a packet number?
Dave B: The MAC gives the PHY lots of stuff (for the PHY to transmit), but this isn't really "giving" it to the PHY, but passing it through the PHY. The PHY doesn't have any business dealing with packet numbers from the PHY.

To close the issue: YEA: 35, NAY: 0, ABS: 0

Report from Group 3 (Jim Schuessler)

Issue 18.5: Is data rate agility only a PHY matter? answer: NO!

WHY?

1) Increase rate in next generation

2) Lower speed to increase reliability (but this may put you lower than 1Mbit, which is viewed as "not good")

Class 1) BSS (other rates in other BSS within ESS) (only works in multichannel environment)

Class 2a) per station (stations are fixed at different rates)

Class 2b) per station (stations are capable of lowest common rate) (you could transfer information at that lowest common denominator)

Class 3) within a frame (Wim contends that this is in level 2, i.e. rate changing is happening and being signaled in the PHY and its header, without the MAC's involvement)

Multicast/Broadcast and how data rate agility may affect these:

All stations must be able to receive multicast without having to send "n" copies where "n" is number of rates. Definition is interoperability (same as ad hoc).

Coexistence

Definition: All stations must support the same medium access rules. i.e. energy detection may solve problem, but length fields break it (everyone must read length) If in the same band and code space "we are all on the same wire."

Ad hoc Networks

Same issues as broadcast/multicast, although bootstrapping it may be harder (due to initialization problems)

Time Bounded services

(Interaction with reservation system) Definition: The one doing the reservation must know the data rate of each station associated with it. The coordination of any speed change and the coordination of Time Bounded reservation need to be one and the same. Only applies AFTER reservation for a connection has been set.

Wim: During a connection, set for a given bit rate, you would not want to change its bit rate (due to effect on time reservation). But multiple connection-pairs can operate at different rates, as long as it isn't changed in any given connection-pair, once the connection has been established.

Power Management

(receive time stamps) Definition: Same as multicast/broadcast -- must be sure all can hear -- time reference. There are some synchronization implications. Negotiated rendezvous are similar to time bounded situation.

Worry of inefficiency of backward compatibility method. (The cost of NOT providing the hooks) Complexity of feature in MAC. (all of the above) Definition: Identify the hook and the cost.

Issues/Problems

Worry about minimal amount of information that needs to be at common speed. Every frame/ "n" frames.

Message to MAC Authors

MAC Authors should address above issues. Assess throughput gain to 2X PHY rate difference.

Assumptions:

- MAC tells PHY which data rate to transmit MSDU
- PHY will tell MAC what speeds it is capable of.
- PHY will tell MAC what speed latest frame was recieived at.
- Matrix of orthogonal axes
- Channelized and non-channelized PHY
- DSSS, FHSS and IR PHY

All .11 MAC proposals. Send message to MAC authors.

Rick White: There is a lot of discussion in the PHY about speeds below 1Mbit/sec. So this group shouldn't automatically ignore rate "reductions" below 1Mbit.
Dave B: Last meeting saw a fierce battle, at which the decision was that we would not go below 1Mbit/sec under any circumstances. This should be a dead issue...

Conclusion

Group agrees that some mechanism for accommodating increased speed migration must be provided in the MAC (PHY also!). Consequences of failing to do it are undesirable proprietary solutions.

Jim Schuessler: We need to attach this conclusion to an issue.
(discussion as to how to get this conclusion into the issues log, or somewhere else -e.g. the "recommendation to MAC authors")
Leon: If the question is "is data rate agility a MAC issue, as well as a PHY issue?", then I think the answer is "yes", and we can close the issue.
Dave B: I think we can do that, and then I think we should open a new issue that asks, "should we *do* this? - i.e. do we want to take on this complexity at all?"
Michael ??: I am concerned with mixing together this issue, which implicitly deals with stepping up to higher rates, and the consequence of sending broadcast/multicast at the common lower rate, and the similar but different issue of "fallback" for increased reliability (being discussed in the PHY subcommittee).

Issue 18.5 Is data rate agility only a PHY matter?

Recommendation: NO

Motion #4 : **to recommend closing the issue as noted above.**

Moved by: Jim Schuessler
Seconded by: Simon Black

Motion Discussion: none

Approved: 36 Opposed: 0 Abstain: 0 ***Motion #4 passes***

New issue: Shall 802.11 MAC support multiple data rate PHYs?

Recommendation Yes, for increasing rates and only for significant throughput gains.

Motion #5: **to have the MAC subcommittee recommend opening this issue.**

Moved by: Jim Schuessler
Seconded by: Tom Baumgartner

Motion Discussion:
unidentified: there's a "con" for identifying it as "increase only". I believe that anyone who adopts this technique would want to include a "fallback" data rate for "increased throughput". If we don't include this, it will come out anyway as a proprietary feature.
Dave B: It may turn out that you have to do that. I don't want to be the guy who has to go out and market it as such: "I'm gonna sell you a Ferrari which will do 180MPH, unless you want to carry a passenger in it, in which case it'll only do 5MPH".
Wim: This MAC is currently targeted for the ISM bands, but it will be targeted for other bands in the future (e.g. 1.9 GHz). The statement that "for increasing rates and only significant throughput gains", this is a limiting statement. We are looking at what things are being done in the MAC for mixed data rates. How much up/down the change is, really isn't an issue in this discussion. We are talking about two different aspects: the mixed data rate situation when you have PHYs capable of falling back for robustness; and step-up due to newer higher-performance PHYs. For instance, the hoppers are working on a 1Mbit PHY, while in parallel, another group is working on higher data rate hopping PHYs.
Mike: We need to make a distinction between units which have multiple-data rates using common signalling (interoperable), and those that do not. The multiple data-rate aspects are very different when you lack the "least common denominator" rate (i.e. interoperability not possible).

Bud: Simple observation in support of getting rid of "significant throughput gain." I could build an 802.11 MAC that only operates at one speed. "Significant" doesn't really add anything - trash it.
Chan: Narrow scope: I would like to keep throughput out of this, though I realize this is the issue Kamilo raised yesterday. But there are other ways to deal with high channel error rates (e.g. FEC). The simplest way to increase throughput is by reducing the reach of the radio link, thereby improving channel reliability. In the narrow sense, I would like to stick with transfer rate on-the-medium, and if the group wants to stick with "increasing", then the reference should be given for the presently proposed PHYs (FHSS, DSSS). If you look at the broader perspective, you should ask if 802.11 wants to address the full range of 802-legal bitrates (i.e. 1-20Mbit/sec)?
Leon: All these discussions are tied tightly with the PHY people. Shouldn't we wait until this afternoon, when they will be in attendance? Otherwise we do the discussion twice.
Tom B: The intent of the words was that it's gonna cost us something to make a MAC that supports multiple speeds. Unless we don't get a lot for that, we don't want to do it. (simple) That's all the word "significant" was included for. Secondly, on fallback, yes indeed this group did not adequately address fallback, so fine we do need to put something in there to allow this particular situation to be addressed.
Marvin: I find the word "significant" fuzzy (ambiguous). May be 10% for one person, 10X for another.
Ron Bjorklund: I view this as a position of direction, and how we should be proceeding. To use the word "how do we do it?" implies that we have a solution, rather than opening a topic for study. I think "significant" is good, because it points to advancement in the future (which we all know will be forthcoming).
Don Johnson: I'm gonna question "throughput". I think of throughput in "Bits/sec/hectare." So data rate alone is not throughput - it's just a part of it.
Ken: Coexistence is a very important part of this issue. Multiple data rate devices should be able to interoperate. Also, by throwing in fallback, this issue is getting very loaded. It's too much for one issue.
Jim Schuessler: Agrees with Ken in separating fallback out of this (new) issue.
Wim: Another strong feeling in the group is that we will need to support backward compatibility.
Simon Black: I think the core issue is just "will we support multiple data rates in the MAC?" All these other points seriously complicate the task of "crafting" and adopting this issue.

Mover and Seconder ask to withdraw the motion. (essentially due to lack of time to continue)

Motion #5 Withdrawn

Issues 15.8 Should all stations support time-bounded services?

Conclusion: Group agrees that "hooks" for time bounded services shall be included in the first release of teh MAC and when fully specified, time bounded services are an option. These hooks are a mechanism whereby the MAC can cause the transfer of isochronous MSDUs in a manner which has an acceptable low probability of collision or deferral. This results in bounded absolute delay and delay variance. The "hook" also includes a MAC field that specifies time bounded or asynchronous data type.

Francois Simon: Functional requirements document asserts that the standard WILL support both asynchronous AND time-bounded service will be included.
Dave B: I want to try a strawpoll. The group forwards the (above) recommendation. How many will support this recommendation?

(strawpoll - no vote count taken. the preponderance

Motion #6: **to close issue 15.8 with the above text.**

Moved by: Jim Schuessler
Seconded by: Tom Baumgartner

Motion Discussion:

Francois Simon: I'm questioning if by adopting this, we are stepping around the functional requirements?

Dave B: I feel that the wording (in the functional spec document) may have not predicted the desire to do this specification in phases (and observing that there is an issue open already to ask if we shall make time-bounded services optional), so I'm not worried about adding this conclusion to the issues log. It will be discussed in the larger group also, so we're not done anyway.

Tom T: I could read this issue to mean that all 802.11 compliant station will include a handset! (obviously pressing the limit of the statement)

Ron Bjorklund: If the question is "should all stations implement TB services" that's one thing, and I would vote NO. But if the question is "should the standard include definition of TB services", I would vote YES. What is the issue asking?

Dave B: I read this conclusion to intend to "leave a hole" for plugging in TB services in the future, so we can move ahead with asynchronous services - NOT the exlusion of TB services for all time.

(discussion of what we are making "optional" here. Too fast to record in detail)

Dave B: I can't get people to the end-point as long as people are driven to vote NAY on anything which isn't perfect. (an expression of the futility of the current train of discussion, methinks)

The conclusion is much broader than the issue. The direct answer to the question (in the issue) is NO (when interpreting this conclusion strictly). The second thing it says is that "hooks" shall be in the first draft of the standard. It doesn't say "at a minimum" (which it should, in my opinion). It also should say that the PDUs should be labeled.

Motion #7: **rewords the conclusion - "Time bounded will be an option, details of how to be worked out."**

Moved by: Dave Bagby
Seconded by: Michael Fishe

Motion Discussion:

Ron Bjorklund: If you go back to functional requirements, it says that all stations shall coexist, whether or not they all support TB services.

Jim Schuessler: I speak against the motion to amend. The group did a lot of useful work. We shouldn't throw that out.

Motion #7 Withdrawn

Motion #6 Withdrawn

General

It seems that issue 10.2 (postponed earlier) is at the crux of the philosophical differences among the MAC proposals. So I propose to leave this issue postponed, and instead, work on the underlying issues.

We need a good foundation from which to do future work. We have 11 proposals on the table. We can't work from 11 foundations to build one house. It seems obvious that we need a single basis to work from. And as soon as I said that, I'm sure a lot of you started thinking, "oh no, how can we do that?"

Nonetheless, I am proposing that by the end of the November meeting, we choose one and only one foundation upon which to do future work. (an allegorical approach to describing an ideal turn of events

follows) I can't control whether the proposal authors work together offline to get us in this position by the November, but I can't see how we're going to make substantial progress until this happens.

Dave asks for feedback. One person applauds.

Simon Black: I'd ask you to put the strawpoll chart back up and I'd like to talk around some possible conclusions from that poll.
Dave B: Should I interpret that as -you are in favor of this, and you want to use the strawpoll as the way to work on this?
Simon Black: I think we should do this, and I think we should work on it THIS meeting.
Dave B: To move in this direction, considering that this is a technical decision, requires a 75% vote. I'll need to get that here (in MAC subcommittee) and then I'll need to get that in the (plenary) group. I realize I may not get there in the November timeframe, but I have to set a goal. I see no more reasonable goal to set. If we set the goal for November, we will move forward. If I set it for 1997, that's the same as no goal (which is pretty much where we are now).
Simon Black: My point here is I agree with what Dave is trying, and this is going to be a *great* deal easier than working with the separate proposals, and I think we should start the work now.
Dave B: I want those who are against this to tell me another way to proceed. I'm willing to listen to that as constructive suggestions.
Chan: I've tried to figure out what the polarizations are that cause these groups. If I try and classify the categories, there is one large group that is strongly preoccupied with CSMA, ad-hoc, or some 802.3 variant, and are not particularly interested in gross capacity. There is another group here, which I think is underrepresented, which is in favor of high capacity and high functionality, and is not pressed to immediately get a product. The former group is pressed to market a product soon. If you put both of those groups in the same meeting, they tend to reach irreconcilable positions. My feeling is that this group has to bifurcate along those lines. I wouldn't prevent or attempt to prevent people who have the CSMA/ad-hoc/low capacity people from progressing, but frankly it is counterproductive to have me in their meetings.
Dave B: My opinion: Such is life. We all know how you get voting rights, how you influence, etc. People are represented to the degree and in the way they want to be represented. If you get 75% agreement on something, then by definition, that's what the "group" wants. I think there are people here who try to represent groups that are underrepresented, but that doesn't add up to votes. I also hear you proposing something like 802.11a and 802.11b; I don't want to go to the executive committee to ask for this. I think this is fundamentally opposed to the objective of the committee.
Kerry: I think the PHY group has already done this kind of bifurcation. As Chan said, there are a large group of members that are anxious to get to market. Perhaps those who are taking the longer view could go off and work on the presumption of, say, the 5.8 GHz band...
Bob: The nice thing about standards is that there are so many to choose from... I would like to point out that we do not serve users. We do not serve our interest by multiplying standards. The value of standards is to *minimize* options, differences, etc. If we don't, we fragment the market and nobody wins. I strongly resist efforts to spin out and go our separate ways. You have to force yourself to come to agree if you want to serve the market, your businesses, and the long-term health of this industry.
Kerry: It occurs to me that we're going to keep coming back to this "deterministic" vs CSMA. It's like putting the token-ring and Ethernet people in the room and telling them they can't come out until they agree.
Dave B: Yes, but their divergence works because they don't mix their products on the same wire. We don't have that choice.
Greg: I disagree with some of the comments that there are some irreconcilable differences between capacity/adhoc and async/time bounded. I believe that we can come up with the standard that can adequately serve the set of requirements to which we aspire.

Ron Bjorklund: I completely agree with that. The differences in the philosophies is based mainly on differences in opinions as to what the marketplace is. If we hurry to draft a standard that doesn't take the "real" market into account, we will serve only a piece of the market.
Tom Baumgartner: There's an equally significant hazard of missing the market by not arriving at a standard.
Dave B: I agree that there are differences in viewpoint about marketplace. What's often not stated precisely is "a single marketplace" or "different marketplaces." Some people are here to do market-building. (They're proceeding to product in advance of the standard). But boy, if I could get interoperability, that would be good for us, and everyone else! That's what many soon-to-be vendors are here for, in my opinion. We spent a considerable amount of time and effort trying to discern the marketplace. But what REALLY matters here in the committee is the 75% vote - irrespective of whether or not it agrees with the reality of the marketplace... I suggest that if you really believe you're right about something, I think you should be motivated to convince as many other (voters) as you can - that's how to get it done. I'm trying to turn up the heat here. It's gotta be done anyway, it's just a question of when.

I want to set the goal that by the end of November meeting to come up with a single basis for proceeding our work:

"Commit to selecting a protocol, to be the foundation used as the basis for future enhancements and refinements, by the end of the November 93 meeting."

Kerry: So if I may interpret such: At the November meeting, we will vote for a TDMA or CSMA-based proposal, and proceed from that.
Dave B: I don't know what criteria will be used to make the decision. This is just one of many issues we must reconcile.
Ron Bjorklund: I would recommend an alternative - for you to request the group for an official response of a direction (process) with which to proceed.
Craig: That's very much what I have been suggesting. It seems like a very abrupt/steep change in the way the work is being progressed. This is not in keeping with a good IEEE procedure. If you chose a process by which to reduce the number of proposals (winnow down) would be reasonable. But to suddenly select ONE out of 11 in this timeframe seems like overreacting.
Dave B: When I took on the job, it was incumbent that I had to arrive at an endpoint. I've been here from the first meeting. It really isn't abrupt - it's a curve, and at some point on that curve, you reach the threshold. We tried in the past to set a deadline for proposals, expecting that many would have come in the last minute with their proposals. It didn't happen. I also see persons who have problems with the extant proposals, but I don't see them bringing solutions to their problems. This week we haven't really made progress on this week's goal (of combining earlier proposals). The problem is that if you don't set a goal, it doesn't move ahead.
Craig: I would like to suggest (not make) a motion. I would suggest that by next meeting, only "combined" proposals will survive further (i.e. you must merge with one or more other authors). It works numerically - it converges within 4 meetings!
Dave Roberts: I strongly disagree that the only thing that can survive is a "combined" proposal. I think it's foolish to say that everyone that has had an idea (on their own) up to now, is somehow wrong. I think they should have the right to move forward.
Dave B: I will add the parenthetical statement "(either existing or a derivative of existing proposals) to my term "protocol".
Bob: This method disenfranchises all those who don't have a proposal on the table. And notwithstanding the view that those without proposals are "parasites," I don't think all of us are. Also, this approach doesn't lend itself to compromise. The things I have seen in other committees that have helped are "comparison criteria" and "evaluations of proposals". Each of us have our own views, but such a method gives airtime to the whole group. Ultimately, it will come down to voting. I will have to look at who I can influence (to vote?) in the direction I

believe should be pursued. Generally applied, this will result in coalitions (outside the committee meetings) which will drive us to convergence.
Jim Schuessler: I understand your frustration, and will probably support what I think Simon is going to recommend (limiting, but not necessarily to one). It's my belief that your frustration really comes from the issues log process, which seems to be unbounded. I think we need a better process.
Dave Roberts: This does not say that compromise cannot happen. It says we're going to get down to brass tacks. I think this will give a basis to work from. Anything we choose is not written in stone - if someone came in with the "miracle" approach, we could embrace it.
Mike Fisher: I agree with Dave Roberts.
Simon Black: I suggest we try winnowing down from the 11 to a lower number.
Dave B: I want to try this, then we can try your approach if I get voted down.
Bill Stevens: We've talked for 2.5 years about the polar issues of "TDMA vs CSMA" and other such issues. We have tried to "converge" (unsuccessfully in my opinion). It is my belief that the fundamental issues *are* fundamental, and the only solution is to make decisions. If anyone has another way to work this, let it be said!

Dave B: YEA 20 NAY 8 ABS: 1

Simon Black: (looking at the strawpoll results) I see four well-supported proposals: CODIAC, IBM, WHAT, WMAC. I surmise from this strawpoll that the other proposals are not sufficiently supported to serve as a basis for the standard. Their contents, however, can serve as input and guidance, nonetheless.

Motion #8: **That we select the following MAC proposals as an initial subset from which a basis for the 802.11 MAC protocol shall be derived.**
- **CODIAC**
- **IBM**
- **WHAT**
- **WMAC**

Motion #8 after friendly ammendment:

That we select the following MAC proposals as an initial subset from which a basis for the 802.11 MAC protocol shall be derived.
- **CODIAC**
- **IBM**
- **National hybrid**
- **WHAT**
- **WMAC**

Moved by: Simon Black
Seconded by: Kerry Lynn

Dave B: (wants to clarify intent before asking for a second) If I were to bring a paper next meeting that was derived from two out of this set and one which is not in this set, would that be a candidate for the basis?
Simon Black: If we have to choose one by the end of next meeting, one from four is certainly easier to achieve than one from eleven.
Dave Roberts: Given the goal we just endorsed, I don't see that this winnowing won't happen anyway, and it doesn't help to do this now.
Jim Schuessler: Since CODIAC and the National proposals are compromise proposals, I believe they should both be included in this list.
Friendly amendment requested - to add the National proposal to the list
Simon Black agrees - amendment accepted.

Motion 8 Discussion:

Dave B: I feel honor-bound to point out that when you look at the strawpoll, you see columns with "lots of color", and others with "very little color". But the range of votes is no more than 4 or 5. Further, it was a small subgroup participating in the strawpoll. So the confidence is not high in the accuracy of this poll.

Kerry: I just want to point out that you said that if you want to influence the process, you should be in the room!

Craig: Two alternatives would be to take that vote in the plenary, and second is to do a written ballot.

Dave B: When I said this is a goal to aim at, I didn't need a plenary vote to "set a goal". This, however, is in a different category. You could argue that this is a technical decision, and therefore requires 75%, both here and in plenary. This is a two-step process to be binding.

Craig: That isn't what I said. What I said is there are two venues in which you could support this action.

Leon: I think this is consistent with us having some commitment to making progress. This says "everybody go home and think about this before the November meeting." I think this is a good starting point.

Bob: I'm ambivalent as to whether to vote this now, or to postpone action to the start of the November meeting.

Dave Roberts: I would like to reiterate that this is superfluous. The previously agreed goal will accomplish the same thing. It's not worth setting an arbitrary cutoff which in essence won't accomplish anything more.

Kerry: This action focuses our work. If we don't focus, we don't get the work done.

Chan: Even given that I am one of the eliminated, it doesn't follow that I won't support this motion. My objection is not to the motion, but I think it should be carried out in the whole 802.11, and I think it shouldn't be left to the end, when attendance has dwindled. We could do this first thing in the November meeting.

Dave B: I'm here until 5pm. If you aren't, I'm not sympathetic. I want to call the question. (no objection voiced from the group)

Approved: 23 Opposed: 5 Abstain: 5 ***Motion #8 passes***

Motion #3 (tabled earlier) is postponed.

Meeting adjourned: 12:11 PM