MAC Subgroup Minutes
Tues. AM, May 12, 1992

Meeting called to order at 8:39, Dave Bagby chairman, Carolyn Heide secretary.

Procedural Introduction by Mr. Chairman

We are not using our time wisely. The process has been he who talks longest and loudest wins, which is not an effective method. Here is a proposal for a different system - which is taken from a paper which was submitted to 802.11 about 18 months ago for submission processing. We are not capturing issues and arguments well at the moment, so we spend a lot of time repeating ourselves.

This procedure says - we will record some issues and their arguments. We will keep a log of what is talked about, what the positions are and what the arguments were that got us to our decision. New people can then review these and come up to speed quickly.

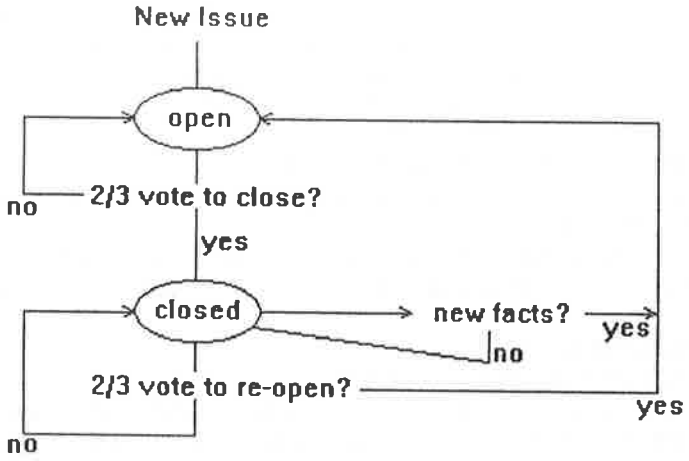

- anyone can open a new issue;

- there will be a log of issues;

- 2/3 vote in support closes the issue;

- issues can be re-opened due to new points on an issue - issues will not be re-opened if there are no new facts. Issues will not be re-opened simply due to a desire to revisit because the conclusion was not liked;

- issues can be re-opened because 2/3 of the people think it should be, regardless of no new facts;

- frame the question; log the issue; move on to other things and revisit to give time to properly consider before voting. If we actually went home and considered issues people who cared could submit position papers and we can all consider these properly before the actual vote. This helps avoid the emotionally of the moment.

Discussion:

Jim Schuessler: In the case where someone wants something and needs approval a 50% vote is more normal. Is 50% still appropriate for motions?

Dave: ANSI uses 2/3 a lot. We can pick any number we want here. Consistency of the percent number for close and reopen is important. If this is a room full of technical experts is a simple

majority really compelling enough? It also makes it harder to get a decision. When we tried to do this in the 802.11 group there was a lot of discussion - in the subgroup maybe we can be more informal.
Vic: Officially for a working group it is 75%
Dave: We require full 802.11 working group ratification to approve anything we do here anyway.
Jim: Thinks this is good , this is really only a subgroup after all.
Dave: MAC has larger attendance than PHY, so we have most of the working group voting percentage in here.So we do have to be a bit careful.
Wim Diepstraten: Should there be a specified time between opening and closing?
Dave: Depends on whether you get the vote quickly. If an issue doesn't get 2/3 it sits there open forever. Play it by ear and see what time works for us. We need to provide enough time so that people who didn't make the meeting get a chance to comment too. Let's try something and give it shot.
Wim: Perhaps at least two meetings.
Dave: Don't wait too long because if we agree we should move on.
Chandos Rypinski: A minimum of one meeting is necessary to close a significant issue so there is a possibility of listening and reacting.
Wim: How about an issue brought up at one meeting cannot be closed until the next?
Dave: Let's try it that way. When people think an issue is going to be closed immediately everyone feels they must immediately repeat their agreement to get on record. This should move us along faster - if you want to say you agree, say just that.
Wim: Should there be a separate record for issues?
Dave: Issues can just be noted significantly in the minutes then extracted later
François Simon: In some working groups it is required that you shall not make a decision without the one meeting delay.
Jim: In SCSI there was a two minute rule - anyone could say when a new issue came up then in two minutes you had to explain what the issue was and everyone had to agree that it was an issue.
Dave: Perhaps I should just decide.
Don Johnson: A paper itself could be an issue.
Dave: We are looking for real issues here - it should be obvious. if it turns out to be problem. We can get more formal then. Asks for an exception to the "keep everything open for one meeting" by making this formal now.

Motion #:1	**That we adopt this procedure for issue approval in the MAC subgroup.**
Moved by:	Jim Schuessler
Seconded by:	Simon Black

Motion Discussion:
François: Do we want the 2/3 vote for approval?
This is generally agreed by the group.

Approved: 19 Opposed: 0 Abstain:2 ***Motion # 1 passes***

General Discussion:
Dave: Everybody may participate and vote in the subgroup. All decisions have to be ratified by the plenary which is the official voting list, so we can do no harm here. Please everyone participate in discussions and votes
Chan: It would be valuable - we have already reached conclusion on other details. People may bring up old issues, perhaps we should find old issues and formalize them now.
Dave: That would a good thing for an off-line group to address.

Chan: For example the 4-step handshake is in many proposals, so that might be one.
Dave: In an evening a small group could do this. An issue should be phrased for a yes or no answer or a very small list of choices. It would be like inventing an issues log and having a quick vote on them. Between now and the next meeting please submit them. We may need to try to inter-relate them. More people may work on things this way by assigning groups to go off line to address issues.
François: I am concerned about where we have been, and I don't know where we are going. Two separate things are happening: people are doing a piece of work, or submitting a proposal for us to look at. We may match the proposal against our criteria - however when matched against the uncompleted functional requirements the proposal doesn't match. At one point we were looking at proposals without functional requirements.
Dave: Hoping this will help. Before we leave a meeting we can say at the next meeting we want to close these particular issues. This gives a way to focus work done in between meetings. People can prepare for specific things that are going to be addressed at the next meeting. In the past, from meeting to meeting we didn't know what we are going to do. People are bringing proposals for total systems and arguing for them to be adopted.
François: These were proposals for designing and we are trying to make a standard, not a design. We were putting the cart before the horse.
Dave: Lets do some issue recording, some arguments, and review new submissions. This meeting will be a bit is a mismash, but you are here to do work - what do you want to accomplish now? papers: 92/49 by Ken Biba (Ken is not here); 92/51 by Wim Diepstraten (estimates 1 hour for presentation); 92/52 by Rajeev Krishnamoorthy (estimate 1/4 hour for presentation); 92/55 by Chandos Rypinski (estimates 1/2 hour without questions). Any other business?
François: How is the MAC group going to work? Right now I don't see the light at the end of the tunnel. We have a bit of a MAC/PHY interface; lots of presentations of MAC proposals which have been kind of matched to set of criteria which don't match the functional requirements as they stand today. Where are we going?
Chan: The functional. req. could be written to define a kind of decision on the system - in order to proceed we are gong to have to have a set of requirements which don't anticipate the answer. Large scale system or autonomous system? If we take a position other than both we can't work - I hope we can define a system that is both - we have got to get off the idea of getting one out by tinkering with the functional requirements. We need an answer - one, the other or both and we better decide both or we are paralyzed.
François: We are dealing with large scale issues which are standard issues. A MAC proposal is just that - not a standard.
Dave: If we had a functional. req. document we could decide what to do to meet it. But we cant afford to sit and wait for that to be done. The issue process will help that - issues will be some small and some broad, some will not be agreed on due to large philosophical questions. This procedure was ripped off the ANSI committee.
John Corey: Is the idea to let individuals make proposals, then we evaluate, then we select? Or is the idea to set down the criteria, then as a group develop a protocol - do we develop, or do we accept someone's development?
Dave: Personally - I think it will be a bit of both. A group will not design a protocol, we will adapt someone's or a combination of several. We will accept someone's general approach and modify it.
John: Evaluation criteria - when a proposal comes in, if a submission complies with our requirements, it would be quicker to decide how much of that protocol is applicable - is it possible for us to set down a list of things that the protocol must do, then submissions can be designed to meet them
Dave: This is difficult, and may be repeating functional requirements work.
John: Not the full requirements, but just the protocol elements, to make sure that it is going in the direction we want to go.

Dave: The thing to do this morning is to review papers. Invites anyone interested this evening to write down existing issues such as - should ad-hoc nets be supported? should infrastructures be supported? (Those are broad, but after those are settled less broad ones will arise) This will begin to constrain us - then when proposals arise we will begin to have a list.
Wim: Isn't this duplicating functional requirements work?
Dave: If it does, then those issues are automatically settled. It would be unfortunate, but we must move on. Basic issues only - we will just capture them, not decide yet. When we get to a written standard there shouldn't be any open issues.
John: You are assuming a single standard?
Chan: There are issues on which agreement is impossible.
Dave: There are a bunch of issues to be answered - for instance do we allow optional features. This would settle one standard or a more flexible one. Let's move on to presentation of papers. Be thinking of issues as we go along. Goals for tonight's issue listing group - not to skip dinner, not to stay up all night, and to make some progress.

IEEE P802.11-92/51 A Wireless MAC Protocol comparison, by Wim Diepstraten

Pages 18 to 40 contain the presentation slides.

Perhaps some common ground for comparison methodology for the future can be found. The purpose here was not to come to a conclusion, just to compare.

Only data is considered, not voice and data (i.e. time bounded).

On page 20, centralized: use of "efficiently" is a personal decision. Will not have enough knowledge to decide on a per station basis which stations can operate in parallel to do data transfer; uses wireless medium bandwidth and backbone (distribution system medium) bandwidth as well for supervisory information; distributed system does not require extra supervisory information on the medium, medium is automatically shared.

Dr. Natarajan asks if this is assuming a single channel PHY? Wim replies that 3 cases are addressed in the paper. Even when multiple channels are available, there will still be some co-channel interference, there is not infinite isolation. Code division isolation has even more obvious isolation limitations. The interference can be significant. But the MAC must be able to operate in a single channel case, so this is clearly the worst case and that where the interference is most obvious. So this is not only applicable for a single channel environment.

When considered here Ken Biba's LBT is slightly modified for 'proper' operation.

Maximum length of 1088 bytes was chosen because this is the Novell maximum (the peer-to-peer and client-to-server evaluation scenarios were also chosen to emulate Novel support).

PHY used is WaveLAN PHY: 2 Mbit/s and characteristics unique to that PHY. Simon Black asks if the PHY includes polarization, diversity and all facilities that it includes? Wim says yes, protocols all analyzed running on the same PHY.

Chandos Rypinski says, the load is generated by 7 stations. Is there no case of 2 stations operating at the same time? Wim: in a single network doing peer-to-peer it is a rare circumstance when 2 in operate in parallel. The model includes that possibility - it will depend on the interference at the intended receiver whether that transfer is successful or not.

This is the same simulator which was introduced last meeting, in submission 92/26.

The theoretical case considered in literature usually has an unlimited data supply where the delay goes to infinite on over-supply. In a realistic model a Poison distribution is normal (and is used here).

Per station, the throughput per station is the only thing that is interesting. When the offered load exceeds capacity it is no longer important. The capacity of a station to generate traffic creates the load, and the limitation is the access delay, which is a function of how busy the network is.

Discussion:
Don Johnson: Actual load (because access delay varies) is just what gets through. You don't loose any packets?
Wim: No, not in protocols that have MAC level recovery. In a high load situation there will be stations in various parts of the delay process.
Chandos Rypinski: There is a load queue behind each station trying to get in .
Wim: There is only a single packet queue in each station.
Don: But more stations make more delay, so 7 stations is a limit.
Wim. Sure, that's why there should be a separate simulation for high population.
Don: This is more realistic than most simulations however.
Wim: The transfer delay is the access delay + transmission delay.
Chan: One packet is about 4-5 msec, there is one contention event per packet, or approximately every 5 msec. So contention events would be larger spaced, depending on the random delay.
Wim: For this simulation only.
Don: This is typical because on Novell there are usually large file transfers broken into 1088 byte packets.
Wim: Goal was to be more realistic than the typical load/throughput literature.
KS Natarajan: What range of delays is used?
Wim: 1 to 320 msec. with a mix of short and long packets because Novell sends short request packets followed by the response with the data. So 60% to 80% of the packets on the medium will be short, so I use a 60% short, 40% long.

The ALOHA curve is straight for throughput versus buffered load because there are NO transfer delays - performance is lousy, delay is perfect. It makes a good reference because everyone is familiar with it.

The delay is only for packets that get through, but there are MAC-level recovery and non-MAC-level recovery protocols evaluated here. So a higher transfer delay reflects the overhead required for the MAC-level recovery protocols. The non-MAC-level recovery protocols don't reflect the lost packet delay.

Discussion:
Simon Black: Of the protocols evaluated 2 have MAC-level recovery, 2 don't - should they not be on the same graph? You can't compare the results.
Wim : That's why I'm explaining it! (appreciative laughter from the crowd)
Simon: I'd like see what happens to the Wavelan system when there is MAC-level recovery.
Wim: The client-server figures reflect this because the recovery is done at a higher level.
Simon: No fragmentation in the MAC is assumed here also?
Wim: Yes. That is why I have graphs for varied packet sizes - I suspect that in a wireless environment we will want to limit the length of the data packets. Large packets are more efficient, but we must compensate for the higher probability of error (especially in the moving station situation).

Note that in the shorter packets, or mixed packet sizes, the curves become less different because the high overhead (for instance of the LBT RTS/CTS) becomes more significant.

Delay is also increased with number of stations due to collision, not just amount of traffic, so there are simulations of performance versus number of stations. You may say, how can there be throughput with 1 station - although we say we have 7 stations, what we have is 7 plus a server, so there is always a destination.

There are lost packets in the LBT - these may be RTS/CTS packets. The point of the protocol is to have these packets lost and not the actual data. Although this may stop lost packets, it may decrease throughput to much. The packet loss rate of a CSMA/CA+ACK has less overhead and not much more packet loss.

Assumption of no error rate, there is always sufficient signal for transmission. The model accounts for interference, not noise. While Simon's simulations focused on noise, I think interference is the limiting criteria. A specified signal-to-noise ratio (SNR) is required for a packet to be successful.

Preamble length required by a PHY is considered - see the graph on page 32. This is very important for higher speed systems where you need, for instance, equalization. Training time is required and it becomes significant. This effect is more significant on high overhead protocols such as the LBT.

Next there are graphs considering throughput versus network separation. The overlapping networks do not show horizontal lines due to: not consistent overlap (some stations overlapping, others not); interference slopes off as the stations get farther away. Chandos adds that eventually it would become flat, for instance one network in New York and one in Amsterdam. Does another station have 0 interference? No, but it becomes weaker. Chan states how greatly pleased he is (as are we all!) to see these different parameters explored realistically.

The area around an LBT conversation which must be quiet is larger than that around a CSMA conversation. However this does not appear to be significant according to the graphs. Wim expected to see this difference.

The large drop at end of bottom graph on page 35 is just an interrupted simulation run.

In the increasing distance between network scenarios, there are effects according to where you are in which network. To create the data, network 1 is moved to the right while the other is held still. So a station on the far left of the stationary network eventually has no media sharing, while that on the left of net 2 has large sharing always.

The graphs which use 20 db between networks are considered to be at inter-floor separation. Chan points out that 0 m difference is unrealistic if 20 db is the difference, to which Wim replies that the distance starts at 5m, not zero - look carefully at the graph, the first point is not exactly on axis!

LBT is more effective on the hidden station problem, so it has less variance in the overlapping network cases. The one network situation is less affected by this problem so the 3 protocols look more similar.

Discussion:

Chan: Hears an inferred conclusion here. Of the two methods - (1) sending whole data packet on meeting of necessary conditions; or (2) experimenting by sending short packets to test the path - you feel that 2 uses more channel time?

Wim: Yes. Throughput of 2 will eventually become better - the fairness of the LBT in the hidden station environment will be better (in this context and for this high load situation - which may be less than a percent of the average load seen during a day).

Chan: Contention does not detract from message transfer only from the capacity of the channel. You have the possibility of loosing some messages while data is being transmitted.

Wim: There is a recovery method available which is only used when data gets lost - that is pretty efficient.

The individual station performance graphs are per 5 second test period.

Submission 91/125 has addition Wavelan statistics. Wavelan delay does not increase due to the p-persistent delay which is invoked during backoff. The standard 1-persistent CSMA/CA is wavelan.

The client to server simulations emulate Novel, where to transfer a packet from server to client, station sends 64 byte request and server returns 512 byte data plus the 64 byte acknowledge overhead. The throughput calculations are based on the 512 bytes actual data. The handshake overhead is not visible on the client-to-server-graphs. In the client-to-server simulation the overhead of the LBT begins to have a good effect. The plain CSMA forces recovery to be done at a higher level, so the MAC-level recovery protocol recovers much more efficiently.

Packet error rate reduces because all traffic is going to the server. 50% of the traffic is coming from the same address.

When increasing the distance between networks in the client to server simulations, distance from the server becomes significant. Also distance between the two servers becomes dramatically important. However the number of stations trying to access the network will be much lower, because of waiting for the server delays.

Discussion:
Jim Schuessler: In the past we have talked about standardizing on a simulator. Is this a candidate?
Simon Black: We might uncover different things by using different tools.
Jim: Yes, but its hard to compare results from different machines.
Wim: In particular the buffered load model should be used.
Jim: Do you consider the code you wrote for this to be public domain?
Wim: I don't want to maintain it - I'll have to think about that.

In the conclusions, when referring to distributed protocols don't get misled by ALOHA! For the distributes systems no additional processing needs to be added for overlapping networks, the medium is shared automatically.

The reason for lost data in the LBT protocol needs to be further explored. Note that if its goal is no lost data, it does not succeed.

Someone asks, does NCR have real data (since Wavelan exists) - comparison to user data would be nice. It might help to judge the realistic-ness of these scenarios. See submission 91/125, replies Wim - this compares real product data. Although the conditions in that data will not be the same as these. A lot of test time would be required to do this in the real environment - I'll think about it.

There was a brief discussion of the logistics of lunch, assuming we reconvene at 2 PM for the Functional Requirements group. This leaves just time for Rajeev to present his submission (since Wim got his presentation in done in slightly over 200% of estimated time).

At 6 PM tonight (or immediately following the afternoon meeting, whichever is later) a small group will meet to assemble existing MAC issues.

Please subscribe to E-mail to allow Dave to easily distribute documents between meetings. Individual mailing and calling are impossible to find the time for. Also remember the CA Microwave number 1-800-2480211 .

IEEE P802.11-92/52 On simulating MAC Protocols , by Rajeev Krishnamoorthy

Since I couldn't do all of the protocols, it seemed that we needed a common framework for comparing results. I like the idea of different simulators, but I think we have to have a common set of parameters for those simulations. Common things like power levels, interference characteristics, etc.

We need to decide upon what applications are we modeling - long to short packet ratios, etc. Ken provided a list of applications, but if we could pick a few, and lay down a basic set.

This paper lists PHY characteristics, transmitter and receiver characteristics, traffic characteristics, auxiliary protocol (like how are we deferring, handoff algorithms, timeouts, adaptive optimizations).

Wim Diepstraten asks what about deferral? **Rajeev:** if you are deferring when the carrier is busy, then sending a packet, or if you defer due to collision - I consider those to be the same deferral.

The parameters of simulation are in the paper. ALOHA is simulated for reference.

We could continue to vary and plot separation, or we could just agree on a set of separations and always plot those.

Wim asks what exactly is meant by LBT in the paper. Rajeev replies LBT exactly - listen, backoff, 1-persistent only. See the paper - there were some changes assumed to Ken's LBT and these are explained.

The simulations tried to varied deferral (strategy explained in the paper). There is the question of what is the proper place to set deferral maximum.

Be careful to watch the notation on the charts, it is misleading, so read the text that goes with it.

Assumptions: throughput does not take into account overhead.

Wim: in general, slot time refers to packet length for ALOHA, but for others it is the total time required to sense a busy network also. Rajeev points out simulations are still in a preliminary stage. He is hoping that before going farther we could agree on what aspects we should be examining, and how we are going to compare results. KS Natarajan asks if channel capacity, should not be a varied parameter, to which Rajeev responds yes - these are exactly the things we have to agree on.

There is some discussion of the possibility of installing a mailing reflector - ie send to one and he forwards to others. Also, on the Internet there are many discussion groups run, we could have one. Although all good ideas, Dave isn't sure that he can find a machine to use for, but he will see what he can do about setting up a reflector, or at very least a group alias.

General Discussion:

Chandos Rypinski: There has been some mention today of acquisition time. All other 802 protocols do acquisition in less than 2 octets. If nothing less is acceptable, the modem designers should fix that. Some phase lock loop circuits require that - the better the SNR the faster you can do that. I expect that we'll do acquisition in 8 bits. We have built crystal controlled 16 mbit/s that do acquisition in less. A long preamble assumption is prejudicial against short message protocols. This is prejudicial against LBT unfairly. Fact - many of these systems have acquisition times of less than 8 bits.

Wim Diepstraten: For cable systems, yes. What I wanted to show was the dependency of a preamble needed by a certain PHY. For equalization on a high speed PHY a long preamble may be needed.

Chan: Equalization does not equal acquisition time.

Wim: But it is part of acquisition time. The MAC level preamble, one of its uses is to cover the acquisition time required by the PHY. When I turn on a transmitter and I want acceptable conditions at the receive end, it will take some time.

Chan: How long?

Wim: That is the training time, the acquisition time.

Chan: Acquisition time, unless there is adaptive acquisition, is going to be under 8 bits.

Wim: Does not agree.

Simon Black: This is not something we can conclude here. The simple point is that preamble length does affect MAC performance. This is an issue which we will need to talk to the PHY group about.

Wim: Yes, but it is an issue we must consider in simulation.

Simon: In the order of 32 bits used in European PCS - it is different as it is a slotted system.

MAC subgroup adjourned at 12:00 noon on the dot, until tomorrow afternoon

Wed. AM, May 13, 1992

A brief MAC group meeting because the Functional Requirements group adjourned early. Called to order by Chairman Dave Bagby at 11:00am, Jim Schuessler secretary.

IEEE P802.11-92/49 Adaptive Distributed and Centralized Coordination, by Ken Biba.

Or "How do small WLANs grow into large WLANs and be good neighbors to other WLANs?".

Presents asynchronous and time bounded delivery services and two data forwarding methods - peer and hierarchical. Therefore, each BSA can have one of four CF - Coordination Functions.

Wim Diepstraten asks for motive behind this. Ken asserts that for market reasons, most near term WLAN requirements can be met by an asynchronous distributed CF, but that perhaps a PCF - Point CF can provide higher network utilization when used with Time Bounded data delivery services. He is proposing that low complexity stations can be invented to provide the asynchronous only service, and with a very small increment, coexist with a PCF providing Time Bounded services.

Stations that don't provide time based services to their users would only have to obey the rules by accepting the beacon MSDU and transmitting at the proper time.

Ken explains how overlapping BSAs, with the same or differing administrations, coexist in both the asynchronous and time bounded, and peer and hierarchical cases. This explanation, of course, uses the MAC protocol Ken has proposed to the committee. This is essentially what the rest of the paper is about. Given this structure, users can now make trade offs of performance verses utilization.

In addition, Ken asserts that this hybrid protocol is robust (performance degrades gracefully) in the face of a wide variety of configurations.

Wed. PM, May 13, 1992

MAC subgroup reconvenes after lunch, Dave Bagby back in the chair, Carolyn Heide secretary (sec: much thanks to Jim for helping out in my absence.)

Something related to,but not exactly, IEEE P802.11-92/55, by Chandos Rypinski

In Chan's dictionary, listen = whatever information you can determine by listening on the channel before you transmit.

The paper 92/55, "Assumptions That Limit Validity in Modeling Listen-Before-Send Access Methods", covers some specific points about propagation.

Take the subject matter and suggest that the generality of Wim's conclusions (in submission 92/51) are much more bounded than were suggested in the paper.

Probability plays a big part. Our par says 99% percent of the places 99% percent of the time - we will find that is recklessly ambitious.

In fading due to multipath propagation there is a limit to how much the path can fade - on the upward side it is bounded, but it can go a long way on the down side. This has to do with the interference of far away stations. If you don't use a Rayleigh curve, it requires justification. There are small probabilities of first a weak desired signal. A coverage area is 95 to 99 percent of the time - an interference range is just there a few percent.

Wim suggests that for the signal attenuation desired the curve should be like that, but for the interference perhaps the curve should be upside down. Yes, says Chan, there are multiple ways of looking at this.

This is where Chan has a problem with the model in 92/51. Two clusters related as function of distance. A cluster size is something about 40 by 80 meters, the distance between clusters is some number and if it far enough away the clusters operate with a capacity of both, but as they get close and overlap they have some degradation. How close can they get and still operate? Each would diminish about 10% - as you bring clusters together they share capacity. He was looking for the point where they drop by 10%. As we push them together we diminish the capacity y of the total. There is a limit as to how close you can push co-channels - this is well known in radio. The important ratio is the coverage radius to the distance between the coverages. A 40m radius at space of 400m is a 10 to 1 ratio. An efficient radio is 4 to 1 with an omni-directional antenna. 4 or 5 times less would be a tremendous loss of capacity. This loss of capacity is just capacity of the station, not stressing the interference factor.

A real situation has many areas used to create complete coverage - how far apart may these areas be run with acceptable diminution of capacity. Every area is surrounded by a ring of stations and users. The interference from the surrounding ring, the aggregate of all interferers, will decrease the capacity more. On a per cluster basis the presence of the ring stations diminishes the capacity of the center station.

What about the space in between the areas? If the co-users have to be some difference apart to operate independently, how do you fill the space between? Channelization - code space, freq, etc. It is not possible to design the MAC without a strategy for reuse of channel over and over again. You have to have a strategy for how to reuse. It is very easy to say the antenna must be omni-direction and all stations the same and get suck with a 25 reuse set. Another thing happens - in the small model presented, the busy situation is clearly defined. But when you get the big picture,and don't have channel reuse (unchannelized) the busy indication becomes the sum of a very large number of systems - you get busy lockout. The range of the system is 4 times greater for busy than for service - or it may be worse. The aggregate of all systems cause the systems to appear to be always busy. If you can never get the channel because its always busy, you have to back off to ALOHA - just ignore what you here! The larger the scale of the system the more possible is busy lockout. There is a point where there is less loss of capacity from this than from interference. There must be a way to include the receive point as well as the transmit point (which is the purpose of the 4-way handshake). The minute these systems become large scale there will be busy lockout.

Access points (AP) and overlapping independent systems have been discussed a lot - Chan suggests that an AP also has a privileged antenna location (on the ceiling). If desk antenna don't stick above the partitions you are going to communicate by reflection with is bouncy and multi-path for instance. If the AP has a privileged antenna, the range is greater - the propagation is better. Nothing makes radio systems work better than putting antenna higher. First you get the improvement of two links with a repeater, but also much lower path loss. The AP will have less rejection due to multipath. No matter how people want to have all antenna alike, the benefits of the privileged antenna AP will be so great they can't be ignored.

It is an advantage to have the propagation poor because propagation increases the interference range. Also the privileged antenna can be controlled because there is only one of it, while there are many others - much greater separation between contiguous systems. A very important point.

Interference between individual stations is much less important than interference between APs - stations have a low duty time, and that will be within the range of the recovery mechanism to deal with.

Discussion:

Wim: What do you consider is causes the diminished capacity? interference? sharing?

Chan: In your case a sharing phenomenon - a capacity division. In a lager model you will also have false busy added to that. An interference limited system is more than just interference. It is one in which the channel is assumed to have signal at all time - you can communicate when the desired signal is above the permanent signal.

Wim: What is "space in between" when you have a completely homogeneous environment - even attenuation as a function of distance in all directions - around the 400 m (the distance considered in Chan's example diagram)?

Chan: Wall loss is an asset in planning the system

Wim: The simulation approach was to simulate a homogeneous environment You have to share that channel over that large an area - that will be a function of the carrier sense level. You indicated that the busy lockout - that level is something you can put a threshold on, so you don't have a fixed relation between them. You can control it and play with it for optimization

Chan: It is a parameter that can be changed for optimization - I don't favour that.

Ken Biba: Diminution of capacity - when two networks are separated you get spatial reuse of bandwidth. As they get close together you get bounded by sharing (if each is 1 meg, separate they are 2. On top of each other they are 1.)

Chan: Eventually you have to share capacity to the point where the capacity is less than one of them operating autonomously. Code space, freq space, time division separation into channels have been discussed - we are going to divide some way, it should be done in the MAC - not by physical means.

Ken: Improving one homogeneous system by adding channelization is the right solution. Make it work for one, then divide it up. Continuous busy - you raise the aggregate noise if you have lots of nodes and create this false busy. We can do more than analog busy, such as preambles for busy digitally detected.

chan: We can improve not fix the situation.

Jim Schuessler: But preamble decreases overall capacity.

Ken: We should build a system that wouldn't like to be channelized.

Wim: What you showed here are the sharing fundamentals for a single channel system - you are right.

Chan: You must accept the capacity degradation or channelize. If we divide the spectrum we are assigned into channels we degrade capacity anyway.

John Eng: Two networks in the same building - the more people (with networks) that move in the worse the performance gets. The building gets full and nothing works.

Chan: The sum of the interlocking coverages needs to be less than the full capacity. It will be hard to get the kind of capacity that video needs - if their traffic is so great that it stresses the capacity they can get more by cooperating (the shopping mall syndrome). 20 Mbit/s to 80 mbit/s per user per hour, but bursty is what Chan has been modeling with - that doesn't even begin to tax the system. The system must be designed for the peak - the nonloaded case is a nonissue. Where are the limits?

Wim: What is the probability of the high load over a certain period? It can be a peak that lasts for a short duration. You can deal with that by accepting reduced per station throughput for that duration. How much will that affect the average throughput of the total system? If you are dealing with that, then the question is, is my system able to work in that overload environment. We have to look at the stability in that situation.

Chan: It has to be stable - you can clip a peak or spread it. The algorithm must decide - whatever the strategy is it should not be random

Rajeev Krishnamoorthy: What do mean by clip and spread?

Chan: Queue of data awaiting service. Waiting time occurs on the queue. Limit is at what point is the waiting time unacceptable. Waiting time access of 2 holding times (ie transmission times) is normal. The waiting time must be predicable. It is possible to specify waiting time in some statistical way.
Wim: Privileged antenna - you say it is important that the AP should have one (ie high up). Doesn't that harm you because your neighbour's AP is high too, so you increase their interference with each other.
Chan: That is important. For a station to successfully receive from an AP requires only that the desired signal overrides the interference by a certain amount. My resolution is that if stations are too close together to operate simultaneously they must operate sequentially. Because of no channelization my system has all the bandwidth, so time sequencing is capacity increasing.
Wim: Without increasing the capacity?
Chan: I'm holding capacity fixed - at least as compared with a channelized system. A case will exist where stations are too close together to transmit simultaneously and that must be dealt with logically. Code space instead of frequency space can be used.
Wim: My simulations do take all this into account. You are now addressing what are the sharing fundamentals.
Chan: We are indebted to you for that work, you stated the assumptions up front which is refreshing.

IEEE P802.11-92/58 Sample MAC group Issue Form and Some current and previous issues as captured from Wim, Dave, Simon, Nat, Jim, François in ad-hoc meeting evening of 5/13/92

A sample form for issues has been created and is the first part of the document François Simon is in charge of the issue log, and the rest of the document contains the issues the group came up with.

Wim Diepstraten: Can there be a hierarchy of issues?
dave: It is almost impossible to sort them such. They are inter-related. Some issues get assigned to groups and the groups get confused.
John Corey: A lot of current submissions bring up issues. One thing that would be useful on a submission, if it would call out specific issues the author feels it raises. The backup to issues may be documents.
Dave Bagby: Organizing these is a manual job. Referring to issues in documents is a great idea. The issue log will go out occasionally. We should give the log a document number, rather than each issue. Issues will be numbered within the log for easy reference.
François Simon: At each meeting we can bring a set of the issues blanks and distribute for people to fill in, like the attendance book.
Wim: On this list are items answered by the functional requirements document.
Dave: We should perfunctorily address them
Wim: We could have a state of adopted, but not closed.
Dave: This evening, think about issues and what to do next meeting. MAC/PHY interface for instance need addressing. Please get papers for us to focus on.
Chandos Rypinski: I brought a revision of the 21 points to the last meeting, that was more structured than the original. I'm not prepared to present that now, but I would like to revisit those 21 points.
Simon Black: There are things in Chan's paper that are issues
Dave: Sounds like there is an issue as to whether we have a set of criteria We need to get some frame we can expand into a standard.
Chan: 802.9 did a great job format-wise. They address questions: voice/data integration, object-oriented management structure, etc. I could edit that down to a skeleton
Dave: Give us that and a justification of why you think its good.
Chan: I will do a form and boiler plate with no content.
Wim: Tools specification (ie for documents)?

Dave: How about some version of MS-word because we seem to be able to exchange these files easily between versions and machine types. Pictures become a problem, but for text this works well.
François: Should we do coordinated work for standard frames, or submit ?
Dave: Whatever happens - we can pull pieces and mix and match if you do things .
Jim Schuessler: There are good ANSI examples too - but we might introduce complexity that we don't need. So use a big knife when you pare these down.
François: Less is better in this case.
Vic: There is also an IEEE submittal standard kit. Look at it to try to keep to it as much as possible. I'll see if copies can be obtained - at least for the editors and subgroup chairs. To order this kit contact:

Michelle Bauer
Staff Assistant
IEEE
908 562-3808 (phone)
908 562-1751 (fax)

Carolyn Heide: About issues - are we just going to have this list sitting around and mark things off as we come to them, or are we actually going to walk down the list addressing each issue.
Dave: We will address them all one by one. The reason I asked for things that should be addressed at the next meeting is that way we can choose which issue to address first.
Jim: At least we should try to group the issues by topic.
Dave: There is a way to specify related topics.
François: Issues will cover several topics too. What should the issue be directed at - the functional requirements, the architecture, the draft standard, the conformance test suite - should issues refer to these?
Dave: I don't know. The PHY group adopt this issue procedure as well, they borrowed the flow chart for it today.

Thurs. AM May 14, 1992

Reconvened at 8:45, Dave Bagby in the chair, Carolyn Heide secretary. Brief interruption from Vic to find out how many people want copies of the venue for the next meeting. Lots of people do, so he will get the copies made. He has already faxed in a request for the standards development kit.

This is a MAC meeting, but the PHY group is mostly are here because they have no room in which to meet. Bruce Tuch is going to give a brief overview from PHY. Take it away Bruce ...

This was a put down everything you think a MAC/PHY interface might need -

signalling classes:

1. data - packet linked
2. control - packet linked
3. management info

each of these has sub-items (see the PHY minutes)

Discussion:
John Corey: This is info that transits the MAC/PHY interface?
Bruce: Yes, bi-directionally.
Nathan Silberman: We listed all the functions we thought needed to be communicated. Then we discussed how that information was going to be transferred.
Chandos Rypinski: Was jabber control seen as a control - isn't that a function that should be entirely within the PHY?

Bruce: It could be, we have to discuss these things.
Nathan: We opted to put everything into the list then delete it later.
Jim Schuessler: Was there discussion of what was in the data band and what was out? Are control and management on a separate channel?
Bruce: We didn't determine that.
Jim: Next you will decide what is real-time?
Bruce: Next is to determine what we really need - if the MAC doesn't need it we won't spend time deciding on how to transfer it.

Bruce puts up an issue list - of issues only relevant to the MAC/PHY interface (see the PHY minutes). Channel is defined as an orthogonal thing - they coexist but don't interfere.

Nathan points out returning number of tries to send a packet may be a good thing to return to the MAC so that the MAC can do something like use the DSS to get assistance from another AP.

John Corey asks if diagnostic issues addressed and Bruce says we didn't discuss that. That may be exclusively a PHY issue and this is interface issues only.

PHY group also adopted the issue procedure. Dave Bagby points out the sample issue form in IEEE 802.11-92/58 . Bruce presented the list of PHY expected contributions for next meeting (see PHY minutes).

Dave suggests that the MAC group should take that issue list to ensure that they get attention

Nathan points out that how are both layers going to communicate with the management layer - the side layer that covers both - is an issue. What are the physical and logical functions that are going to go there?

IEEE 802.11-92/58 - issues list section

Note that the numbers on this list are just to make discussion easier, they are not issue numbers. As only about half the group has read it, we will pause for a few minutes for people to read it and jot down anything that it brings up.

New Issues

1) Channels (PHY group)

(1) same channel/AP

(2) different channel/AP

(3) both of 1&2

2) if 1.2/1.3/1.4 supported then do we support seamless handover? (PHY group)

3) define seamless (PHY group)

4) is mac support needed for antenna diversity (PHY group)

5) is mac support needed for power control (PHY group)

6) does the phy perform or support the security function (PHY group)

7) do we supply a phy type to the mac (PHY group)

8) is mac/phy exchange needed to supply network management information (PHY group)

9) does the mac supply a packet number to the phy (PHY group)

10) is data rate agility only a phy matter (PHY group)

11) what are the environments including station speed (PHY group)

12) what is the intelligence level of the mac/phy interface - same as #51. (Dave Bagby)

13) is the layer that provides phy independence the same as the mac/phy interface or not. (Dave)

Discussion:
Chandos Rypinski: Assume that it is by default - "where are the problems in doing that" should be the issue. Assume so and look for problems.
Dave: Don't assume the answer in phrasing question.
John Corey: Can there be PHY independence, allowing for multi-phy, without affecting the interface. Could I switch PHYs without telling the mac?
Dave: That is included in this issue.

14) what the dss functions needed. (Dave)

Discussion:
John: in #42 (document 92/58) does "internals of the dss" means functions?
Dave: I think that means how does the dss work. Let's not spend too much time getting the wording straight.
Jim: But we need to make sure everyone understands the points.
Dave: Could the issue log editor sort the issues out so that there is less overlap - we give him the licence to rephrase and combine?
François: yes.
[The group also agrees to give him that licence.]

15) what mac users other than LLC do we want to support. (John Corey)

16) what are the logical and physical functions required to communicate to the management layer. nat what is the relationship between mac phy and network management? (?)

17) what are the trades offs in efficiency between a connection oriented protocol versus running time bounded data over a connectionless protocol. (?)

18) what kind of services will be supported with the time bounded services. (Don Johnson)

19) where shall the connection and connectionless services by integrated - the mac or llc or somewhere else. (Chandos Rypinski)

20) what is the algorithm for managing the partitioning of capacity between the time bounded and asynchronous services. (Chan)

21) what is roaming? what is handoff - same as #12. (John)

22) do all stations and all infrastructures support the time bounded service. (Simon Black)

23) what kind of error recovery mechanisms are to incorporated into the mac. (?)

24) do we intend to define or limit the max number of stations in a BSS due to media characteristics. (Nathan Silberman)

25) is the ds interface at the mac or the phy. or both (John)

26) what is the strategy for capacity control. (Wim Diepstraten)

27) what is the mac frame structure. (Jim Schuessler)

28) what are the performance requirements of the ds. (John)

Moving off new issues, let's determine what do we do at the July meeting? This will help us focus on some specific issues. Mr. chairmans suggestions:

1. Issue log flesh out: add alternatives; add pro/con arguments.

2. Define MAC/PHY interface.

3. What are the distribution system services.

John ?: sequence of issues - should they be ordered somehow?
François: A lot of issues are related/duplicated. They will be concatenated, or listed together at least. Seems like now we are doing MAC/PHY interface - are we going to do this together or is the PHY doing it too?
Dave: Lets keep separate MAC and PHY issue lists. If issues are on both, who cares for now.
Simon: We have a good handle on what async service means and its implications. We don't have a good handle on the time bounded service and the services we need to provide have a large impact on the things you have said we need to discuss next. So we need to address this early. It is a substantial issue that should be handled early on.
John Corey: There are a number of very large issues like this that we have to address first or we could waste a lot of time.

4. What does time bounded mean (#9 from 92/58)

John Eng: This goes back to the functional requirements. Without those requirements tied down we can't do the architecture.
John Corey: Back to allowing options in the standard - how you let cost (dollars, spectrum, etc) effect your decisions. This is one of those things that must be settled early.

5. Set policy regarding options in the standard.

Nathan: Cost is a four letter word that may not be discussed, historically in 802.
John Corey: Budget would be a better word - not just dollars - spectrum, power consumption. What are the trade offs in evaluating the choices I want to make. If we make an optional system people can build conformant equipment that implements a subset and can be more portable or whatever. Maybe we need a call for a submissions that describes how to do the trade offs. IE the semiconductor people may be able to tell us, to get all the functions into a chip you can or cannot do it. We need information - not answers.
Dave: You are looking from viewpoint of the person answering the question.
John: Just resolving the options policy will help.
Simon: What is the conformance level is another way to phrase this options question. Are there classes of conformance? We have to write a conformance test standard at some point.
François: According to the international standard (ISO) for conformance - if you state that your product has an option, your product must be tested according to the standard. If you implement that option you must adhere to the standard - the entire option.
Bruce: Is a conformance class an option?
Dave: Modem analogy - if there was a modem standard (ha,ha). Supposing the options were 300, 1200 or 9600 bit/s. You could say you are conformant if you had the 1200 option - you cannot say and I have two out the 3 things required to do 9600 so I am conformant with that too.

Attempting to find out which of the 5 items listed above for consideration is most important, we vote.

On which of 2,4,5 is the most important (item - vote count): 2 -12; 3 - 4; 4 - 0; 5 -6 .

Carolyn Heide objects that 5 is not a MAC group issue, it is a plenary issue. Dave responds ok, so we will go to Vic and insist that the plenary meets on it.

On the second most 2nd most important (item - vote) : 2 - 6; 3 - 1; 4 - 4; 5 - 10 .

So, we had better so 2 and 5.

How many people think that next important is 3 over 4 (11). The other way around (10)

Discussion:
Nathan: Since we know what we are going to address, can we decide what we should have ready to discuss it?

Dave: Whatever we decide is our most important subject we will ask for submissions on that subject and consider them above any others.
John ?: How can we handle any one of these issues without the others?
Bruce: It does have to be done in detail.
Don Johnson: First you have to decide whether the interface is a boundary. Make a decision and change it only if new information.
Dave: We are not looking for details, just generic decisions and a lot of questions. Like - are we going to move data across this interface.
John ?: Isn't this more a subject for subgroups. As opposed to 5, which is more of a full working group argument subject.
Dave: We are going to make sure 5 gets handled in a plenary, early. Within the MAC group we will talk MAC/PHY interface. And if there are some people who want to do 4, they can do it. That's how the agenda will be structured. So - look at the issues list, think of issues on these subjects.

There is some discussion of the relative impact of 3 and 4. The emphasis is on what affects the MAC/PHY interface. What issues affect the MAC/PHY: 1,2,3, 6(?), 8, 9(?), ... in trying to continue we find there is some serious discussion about whether each item affects it.

It seems difficult to determine what issues are relevant to the categories to be addresses next meeting. Maybe the right thing is to let people decide for themselves and submit things that are relevant according to their own judgement. If you think an issue is relevant to the interface, say so in a paper. There should be some papers defining what intelligence lives on which side of the interface.

Vic wants to know if there will a Monday morning meeting in July? Overwhelming consensus is no because of the long weekend.

Possible contributions for next meeting:

- Jim Schuessler: How we accommodate both smart and dumb phys. There is a market for both types of products which are 802.11 conformant. How the interface should allow provision for both.
- François Simon: Model of MAC/PHY interface architecture, based on the OSI model.
- Nathan Silberman: Options and choices for the MAC/PHY interfaces.
- Wim Diepstraten: Functional requirements of MAC/PHY interface
- John Corey: Appropriate description of handoff (functional requirements submission)
- Jim (for National Semiconductor): Something regarding time bounded definitions.

General Discussion
Wim: Do we have enough feeling for the components and detail of 802.11 function to make the decision about options in the standard?
Dave: I don't know. Since next meeting is plenary there will be a lot more discussion too, new opinions.
John: Are we going to limit the number of PHYs, or will we continue to accept new PHY types as people come up with them?
Dave: In the past there has been resistance to cut-off dates. One of the design criteria is a MAC that deals with multiple PHYs, so new PHYs can be adopted.
Don Johnson: 802.3 adopted a set of criteria you had to go through before you could submit a new PHY.

MAC subgroup adjourned at 11:20 AM.