IR-PHY Adhoc Group Meeting
Monday, July 12, 1993

The meeting was called to order by Tom Baumgartner 8:30 AM. Carolyn Heide secretary. Agenda is document IEEE P802.11-93/93.

1. Adoption of Agenda (P802.11-93/93)- by consensus

Larry suggests addition of item 7.1, schedule for first draft of IR PHY spec.

2. Comments on and corrections of minutes of the last meeting (P802.11-93/91) - none, approved by consensus.

3. Presentation(s) on methods of conformance testing of IR transmitters - none.

Tom asks if anyone is aware of existing procedures that might apply. No-one.

4. Presentation(s) of "straw man" specification of IR PHY - none

5. Presentation(s) of other submissions - none

6. Plans for preparation of a submission to MAC group regarding how the MAC can affect the IR PHY

Tom thinks MAC group will get seriously down to specification next meeting since this is the last time they will accept new MAC proposals, so we need to get our input to them.

Tom suggests timing may be one issue - FH has a when to hop timing consideration, for example. Does IR has an equivalent? All PHY's will need to tell the MAC how long they take for warm up, or stabilization, so is this a parameter that gets passed to MAC?

Rifaat Dayem: has another PHY group generated a document that can be used as a template?
Larry van der Jagt: the template which started from Nathan Silberman's document has a column for DS and another for FH (93/83 from the May meeting). Rich Lee also did one for IR a long time ago. A move toward a template like that would be good. We have told the MAC that single channel support is a requirement.
Tom: applicable items from 93/83 are: minimum number of channels; transmitted power levels ...
Larry: do you think you will specify a minimum or maximum?
Roger Samdahl: both.
Tom: set a power level then individual products are only allowed to vary by x%. Sees giving at least a hi power range for normal operation and a low transmit power range if you really want to save battery power.
Rifaat: if you have a lot of people in a cell you may want to turn the power level down - to avoid the near/far problem.
Tom: I'm biased because in my vision of a MAC the near/far problem doesn't exist. Or antenna need to be set-up so that near/far is not a problem.

Roger: don't assume IR will fill up whatever cavity you are in. You may want separate cells in a space. Asymmetry is a problem, where your transmit and receive range are different. Simultaneous transactions in the same physical area and minimum and maximum power level setting in a single area are two issues.
Larry: Rich Lee's 1992 paper had a diffuse and directed column. Issues are going to be different. Need to decide which or both are going to be standardized - directed, diffuse, or both.
Tom: proposes only diffuse in the first version of the standard.
Edwin Turner: what about point to point links, such as docking stations that are very directed?
Tom: There is a big application for moving portable devices. Diffuse serves the need of point to point as well as that need, but directed only serves the point to point need. That is why I think we should concentrate on diffuse.
Edwin: LAN bridge approach as extensions exist already.
Rifaat: such as Infralink.
Larry: the PAR says mobile is the application.
Carolyn: if people want to see directed addressed here all they need to do is bring submissions.

Straw poll: initial focus on diffuse-14, directed-2, abstain-9.

Tom: OK focus will be on diffuse. Back to power levels. Sees two - one for battery conservation, and one for filling up as much space as possible. With diffuse IR the power level specification number depends on the method of the power measurement.
KC Chen: transmit power depends on data speed - without specification of data rate it is meaningless.
Roger: modulation scheme is important too.
Leon: the point is we want to tell MAC group that there are at least two power levels, or more.
KC: should we try to make the same power level numbers as the radio group?
Tom: that would be just a coincidental if they were the same.
Larry: if the top bound is the same that would be good. The number of levels for which bits have to be reserved is what's important. DS is talking 7 levels, FH 4, you are talking 2, so that's OK. Number of discrete levels is all they need to know
Larry: adjustable data rate is another issue.
Tom: a couple of meetings ago I felt that for the first specification 1 Mbit would be appropriate for the standard. However my company says I should suggest 4 Mbit/s. The reason we picked 4 Mbit/s for the Spectrix implementation is because with the components the way they are, 4 is the maximum achievable practical rate these days.
James Harrer: how about 1 as minimum?
Tom: for interoperability some rate has to set. Products from two different companies must talk - either standard rate or something negotiated.
Roger: would still like to see 1 Mbit.
Tom: what about two rates and some negotiation technique?
Roger: we are bound to 1 at low end by the PAR. We would shot ourselves in the foot by not defining some way to negotiate higher for the future.
Henry Ngai: wireless will be a cell of a LAN. In MAC layer they have some provision for support of voice/video, we should not put ourselves in a position of stuck to 1 Mbit which doesn't support the 1.5 for video. Thinks the minimum should be 2.
Roger: there will be applications that are not video dependent and battery power is a driving factor, so some versions don't need to be 1.5 Mbit/s. A negotiation is the right approach.
Richard Ely: video can even be made to run slower.
Leon: what we are asking the MAC is how many bits are needed to describe this number of data rates for the negotiation required. Other PHYs have the same sort of things - certain discrete rates. One is the standard all PHYs must handle, and they can support others from there.
Tom: can see 1, 2, and 4 Mbit, and knows some people would like to discuss 10. There are at least 2 bits worth of choices.
Larry: just the fact that a set is needed is what's important now.

François Le Maut: if you want several co-located cells 4 Mbit/s will be difficult. You would have to channelize.
Edwin: what about very short range, very high speed - a docking station type. In IR trading speed for distance is much more important.
Tom: we defined three important applications in the first meeting and one was docking station. Keep raising those directed IR issues so we don't design them out.
Tom: input receiver sensitivity - is that a MAC concern item?
Larry: anyone got any numbers for this?
Tom: the first question is what are the units - perhaps mW per sq cm. Not sure what the number should be. Will propose one for next meeting.
Larry: sure it's not microwatts per sq cm?
Tom: maybe.
Larry: two numbers involved - the minimum and maximum at which you can make the required bit error rate.
Tom: dynamic range has to be a specification too. I believe you have to walk right up the transmitter and have it work. It is a very wide dynamic range.
Roger: input sensitivity might be controlled by MAC?
Larry: sure - you're getting bit errors, so MAC can increase sensitivity or decrease sensitivity and see if it gets any better.
Leon: MAC may have to adjust transmit power too. Maybe the receiver tells the MAC what the receiver level is.
James: variable gain devices in the receiver under control of the MAC.
Tom: yes, sees adjusting receiver not the transmitter since many receivers at various distances receiving one transmitter. Adjusting receive on the master in master/slave protocol according to what station you are talking to.
James: an AGC situation where you vary the gain depending on the input level.
Tom: implementer has choice of putting smarts in master or slave - others have peer-to-peer architecture in mind too.
James: have to have the freedom to allow it to work in any architecture.
David Fisher: think of problems in analog phone systems. In order to keep error rate down, the power level is arbitrated frequently. You can reduce the recurring cost of the slave by putting the smarts in the master only. A link for the MAC layer where the S/W gets some information based on how well packets get received.
Larry: don't make gross assumptions about the MAC - telling the MAC they must operate in a centralized mode. If a point co-ordination function (PCF) is required for IR PHY we better tell them now.
Tom: from the radio PHY meetings - has this been an issue?
Larry: we talk about it all the time. I think you can say that the MAC group has decided that you have to be able to operate with a distributed co-ordination function (DCF). [Carolyn agrees with that conclusion.]
Tom: would be disappointed if we had to force a PCF for IR.
Rifaat: but you get a DCF by having one of the ad hoc stations becoming the master.
Carolyn: that is a distributed PCF.

Tom: preamble length - do we have to tell the MAC what it is?
Larry: if we use the primitives we discussed a long time ago for the MAC/PHY interface we don't. There was a start of activity versus a start of data. But - how long the preamble has to be is important. The length of the preamble may establish the throughput of the system because there may be many small control packets. If you could come up with an idea of how much you thought, it would be useful. DS says 146 symbols, FH is at 32 bit times.
Tom: in the Spectrix design 3 octets is the choice - 24 bits regardless of data rate. Any other numbers? [no one speaks up]
Edwin: any reason why different PHYs need to be different preamble lengths?

Larry: radio might be making diversity selections in the preamble which makes it longer. IR may not be planning to do this.
Edwin: must be a benefit from higher levels - efficiency gain?
Carolyn: some MACs have to know the preamble length because it affects the length of things like fixed length time slots.
Larry: different PHYs will have different lengths. Say what you want now so that it gets considered if we ever try to set a standard one.
François: you determine speed based on the clock itself?
Tom: hadn't thought about negotiation of speed. Methods for negotiation of speed haven't been discussed.
Larry: thinks PHYs will decide on minimum set of services. All PHYs must operate at this minimum. If they have ability to go higher some higher level entity above 802 will command the change.
François: but in the analogy to modems, control is in the preamble to make the selection at the PHY level.
Roger: it is possible to have a common preamble that is sent regardless of speed. Two units which can't quite communicate, but can detect each others presence from preamble and operate with politeness toward each other even if they can't communicate.
Larry: so for a 4 Mbit system the preamble has frequency content within the 1 Mbit band?
Roger: do your clock recovery in the preamble that was originally designed for 1 Mbit.
Henry: so MAC could switch between 2 speeds, and stations change to talk to other stations. But that brings everyone to the lowest common denominator which is not fair. The preamble and the MAC should work together to allow negotiation of speed using the preamble. The RTS/CTS exchange in the MAC could be used to negotiate speed.
David: preamble specific to data rate?
Henry: MAC already has RTS/CTS that could be used.
Carolyn: be aware that not all MAC proposals do this exchange.
Henry: just using that as an example. The basic point is don't degrade the whole system for lowest station.
Larry: would choose to make everyone in a cell degrade to avoid the complexity of multiple speeds. NCR has brought analysis of different channels used for different speeds. Media arbitration used. Everyone at the same rate is better.
Henry: about the modem analogy, remember they are point to point and what they do affects no one else but those two conversing.
Edwin: Ethernet says everyone works at the same speed like it or not.
Tom: and they added two more media and maintained that too.
Carolyn: but people don't accidentally wander into your Ethernet LAN. Your Ethernet LAN is not affected by the guy across the road.
Carolyn: please bring it up to MAC group - they are wearing blinders on the issue of multiple speed support.
Jim McDonald: can see the gear shift on data rate, but it's logical to keep exactly the same structure just slow it down. For instance would there be different header sizes depending on speed? Does data rate change imply just speed up or slow down, or different structure?
Tom: sees it as just speed, not structure.
Larry: preamble may be affected. If a certain number of transitions needed OK, but if certain amount of time is needed that is different.
Tom: document 93/83 discusses paths that go between PHY and MAC. Two data and a clock - common to both of RF PHYs. Then the lines diverge between the two RF PHYs. Let's take the FH list and see if it satisfies us. Two data, clock, valid data (Larry says maybe that means bit sync acquired). Anyone think bit sync needed? [no answer] What does the MAC do with that information?
Larry: important to know in frame or out of frame. So a frame delimiter of some kind is needed. MAC needs to know when to start looking at what to include in the CRC. Maybe this isn't the

time, but how to determine in/out of frame is needed. If MAC is not going to decode frame delimiter. But this list you are looking at is the physical lines totally within the PHY because the exposed interface is not at the MAC/PHY interface - it is within the PHY. These are things that go between one piece of us and another piece of us if the interface is exposed.
Tom: this document says convergence layer and I was assuming that was the bottom of the MAC, and it is the top of PHY.
Larry: when we get farther along a subgroup of people is going to define that interface if exposed. Will look at existing ones and think about what the physical connector wants to be and how it would be used. Until then it's irrelevant.

Tom: reads list from document: signal quality, loopback, # of channels available, channels in use ... everything we have mentioned is there except preamble length.
Carolyn: what about wakeup latency. After I give PHY the wakeup command there is certain amount of time before H/W actually comes awake and MAC timing needs to consider that.
Henry: what does wakeup means?
Tom: if implementation has sleep mode for battery saving there is some stimulus that says go back to full operation. That may take some time.
Henry: higher layers - from time to time heartbeat is sent. Do I have to wake up for this?
Tom: NCR has a paper that talks about some of this. Implies the AP would answer that question for stations that are asleep. A spoofing type of operation. Battery conservation conflicts with heartbeat operation. Has heard somewhere that Novell is working on mobile support and hope this is one of the issues addressed.

Roger: which direction (MAC to PHY, or PHY to MAC) is important to add too.
Tom: number of levels of receiver strength from PHY to MAC.
François: disagrees.
Larry: PHY tells MAC levels I am going to report at are this, this and this.
James: in currently implemented analog systems there is a continuum of levels, monatomic.

Henry: indicator from PHY to MAC saying wake up now. Such as because data received - something happened on the network the MAC should know about.
Carolyn: troubled by that - all indicators that the MAC might need to know about are there. If PHY gives one or'ed combination of these to wake up a sleeping MAC that is an implementation choice, not a specified requirement.
Larry: other things might happen in the world that the MAC might want to know about after the MAC has decided on its own to go to sleep.
Carolyn: then there should be indicators in the list for those things individually. Its no good waking the MAC up if you can't tell it what happened.
François: at the MAC level this intelligence needs to be there so this additional function is not needed. PHY doesn't have intelligence to know what the settings are.
Larry: we are talking about primitives, functions, not implementation lines. PHY management primitives with bits set in them - standard ISO stuff.

Final List MAC/PHY interface information

Transmit power level >= 2 (no more than 3 bits) (bi-directional)

Data rate: 1 Mbit/s, 4 Mbit/s (bi-directional)

Receiver sensitivity selection (bi-directional)

Receiver signal strength (bi-directional)

Preamble length (to MAC)

H/W wake up latency (to MAC)

7. Plans for future actions and meetings

7.1 Schedule

Larry describes the overall PHY layer schedule known so far:

In FH PHY there are 3 open issues (1) what modulation method; (2) what needs to go into the air to achieve synchronization and hop acquisition; (3) what is the point of conformance for FH PHY in air or on cable. They will close the first and work on second at this meeting. In September they will close the second issue and work on what goes into draft standard. In November we will see first working draft.

A schedule has been requested from DS.

Tom: must shot for an IR schedule by November.
Larry: a draft standard first presented in November (work on it in Sept. and try to get someone to do the work by Nov.)

Start at the end and working forward results in:

November - presentation of draft.

September - transmit levels, receiver sensitivity, dynamic range, testing methods.

July - modulation; preamble

Tom: Spectrix would like to get this done, but won't have time to contribute much in this time period.
Roger: this schedule is right. Not sure that Photonics would be able to define what were doing as a driving force.
Larry: we need an editor volunteer. Someone to take output as it stands in the meetings and other standards and make a template. After that the document is not the editor's anymore but gets changed as input and arguments occur. The other groups (DS and FH) don't have editors yet.
Richard Ely: how much work is that?
Larry: 3 weeks to a man-month to get first draft in shape. After that probably a week after each meeting. It's not a small effort, especially in the beginning.
Tom: Rich Lee's document gives a good idea. Would the 802.4 fiber optics specification be a good starting point?
Larry: yes. If Manchester or BPFSK chosen for modulation then there is something there already.
Rich: will ask management about doing editor work, but is not very optimistic.
Rifaat: adding a column in that table talked about today (from 93/83) would be a good first start.
Tom: what are the chances of people coming to the Sept. meeting prepared to put input into the draft specification? None. If we get input on those items in Sept. and an editor we may be able to meet the schedule, if not we won't.

Modulation Method

Possibilities: Manchester, PPM, BPSK, BFSK, RZBI

KC: what kind of PPM? 4 array or 8 array may effect whether people choose PPM. 4, 8 16 are not equivalent.
Tom: only 16 has been discussed so far.

Larry: Manchester, BFSK, BPSK are kind of a family. Manchester is BPSK with carrier frequency = symbol frequency. The two FSK and PSK including Manchester are things that give opportunity to operate in the baseband (or as close as you come with Manchester), or shift the

carrier frequency and offer ourselves reuse and other neat things. These schemes open a window of opportunity that we should leave open. Not close the door on reuse, co-location.
David: PPM and RZBI lock you onto baseband and that's why they should be eliminated?
Larry: concerned about PPM in the real world. Don't see what you gain by RZBI, except maybe cheaper. The other things are produced in component mass production. A high degree of commonalty with things in the radio world - a whole wealth of things.
Roger: not sure can talk for PPM, but we (Photonics) are use 1 of 16 PPM. The original paper of Dick Allen's is a review of the Photonics system. Simple to implement. Effective in terms of power utilization with a given bit error rate.
David Fisher: what about multipath?
Roger: in ranges we are concerned with it's not a problem. 1 Mbit, 30 feet - the propagation velocity doesn't make multipath a concern. It will be at higher rates and larger areas.
KC: problem with PSK. Everyone will probably use cheap LEDs, and the phase noise destroys the performance. Need to look into the phase noise impact on bit error rate performance. Phase noise = we do not have a coherent source so we are trying to use non-coherent and build waveforms to build BPSK.
Larry: doing electrically and using optical as carrier. Non-coherent at optical carrier frequency.
KC: to not consider phase noise impact on electrical waveform is dangerous.
Tom: a study of whether that phase noise problem is relevant to bit time at these rates would be interesting.
KC: need to make decisions not on feelings - this is science not religion. Non-coherent radiation, so using optical energy through detector to build waveform doesn't get a very precise curve. Not a trivial system to do this. Propose some kind of study needed.
Larry: if you can do Manchester you can do BPSK.
Tom: doing 1 Mbit in Manchester, the LED would be on or off in the order of a microsecond (a full bit period). Any phase jitter optical problems have to be dealt with. In BPSK what is the maximum time light is on?
Leon: one cycle of whatever the carrier is.
Larry: Manchester is 1 MHz carrier. Move up a couple of KHz to get out of common noise sources.
Tom: RZBI - real simple, real cheap, and still works. Very good power conservation.
James: what level of bit insertion do you propose?
Tom: often enough that power spectrum stays above 300 - 400 kHz. At 1 Mbit that's a lot. At 4 it's not so bad. It would seem that RZBI is not good for 1 Mbit.

Tom: we've got two companies producing products with 2 methods (RZBI and PPM), and a lab produced system working with another (BPSK), and another in university (Manchester)!
Carolyn: the winner could be the only one to bring a submission next time.
Tom: there may be issues people don't want to discuss because they're too proprietary.
Leon: make a matrix with some of these things across the top - noise tolerance, channelling, cost etc. - helps decide whether they fit and how important they are.
Roger: this table should constitute things we consider important enough to make the decision.
James: why can't different modulation schemes coexist and be compatible and have different systems recognize others? If it's a wash for instance.
Tom: thinks the answer is that one of the jobs of the specification is that if different company's units walk into the room they must talk. So if there are multiple modulations schemes you still have to fall back to the common scheme.
Larry: using 802.11 as a meet me channel to decide on operating characteristics is fine.
Rifaat: need for channelization in IR versus radio? I look at IR as a low cost alternative and am wondering if it can be channelized?
Larry: doesn't believe there's a big cost. If you look at Manchester as BPSK then you have filters, no mixer to bring down to baseband. Doesn't see the cost. You can get channelization.

Tom: if I was convinced of that I won't have any problem. The need for channelization - Larry has identified it. The cost is another story.
Larry: is there a requirement for uncoordinated co-located operation? If so, then there is need for it. In a big floor there will always be a point of overlap, but does there have to be co-ordination there?
Tom: slightly different question - is it more cost effective to find a means for co-ordinating a large area or to find a means to have uncoordinated co-located? Has already proposed a method of covering a 32,000 foot area.
Larry: because of the limited propagation.
Tom: application to be served is a key.

Here there is an inconclusive discussion of adhoc'ness and what it means.

Jim: channelizing IR - what are you envisioning?
Larry: flashing LEDs at different rates, just like picking different frequencies on a radio. Intensity modulating the light. Envelop of light becomes the waveform and then you FM detect. Centered at 1.5 MHz, 3.5, and 5.5. Spectrum is almost free so we can be pretty gross. Use same LEDs and photo detectors - but don't have to be careful about our side bands.
KC: and then you have fading.
Jim: send out a much higher rate of pulse and turn it on and off at some duty cycle. To have separable channels ...
Larry: channel is not any wider than radio. FSK or PSK them so that they occupy a band that is 1 MHz wide.
Jim: don't you lose a pulse, you get obliterated by someone else's pulse?
Roger: have to have a linear receiver, not just light on or off, to determine signal. Have two LEDs on at one time and be able to separate in the receive circuit. That's where the cost is. You go digital immediately after the pin diode and you can't recover this.
Larry: - you're filtering anyway, just use a different filter. Discriminator of the frequency is so low - for same price center at 10.7 MHz shift from 20 down so that for 37 cents buy the same filter used in radios.
Tom: there are applications where channels would be nice, but has a biased view that the applications we should be addressing are not those where multi-channels should be required.
David: multi-channel where single channel works within it?
Tom: that choice would be Manchester.
Leon: let's get back to listing criteria - multi-channel support or not is one of the criteria. People need to think about these things and bring arguments in favor or against.
Richard: we need some data about the pulse spreading in the room. Don't get pecular reflection you do in radio.
Tom: the papers from Rui Valadas last time hint at that without directly addressing it.
Richard: you send out a pulse and what comes back there are spikes. Not as bad as in radio.
KC: thinks there is a paper in the Nov. 91 issue of IEEE Network magazine (special issue on wireless networking) by John Barry from UC Berkeley that has exactly what you want. Simulation work. Ideal optical path around 40 to 50 nseconds.
Maurice France: multiple data rates requirement - compatibility between channelization and multi-data rates envisioned?

Leon: multipath sensitivity, affecting maximum rate and maximum distance.
Roger: power consumption.
KC: combine SNR with power consumption - given this SNR how much power do you need.
Tom: SNR is "given all transmission at the same power level what is the SNR versus BER". Power consumption is if this technique to transmit x bits, some other technique takes a different amount of power.
KC: depends on BER

François: and speed.
KC: if you draw the SNR curve and BER curve for each technique it will be very clear which one is better.
Leon: peak power versus average power? Short high power pulses versus wide pulse with average power. These affect things like diode life.
Larry: high peak to average is good for life.
Tom: real high is not good - doesn't heat up, but does something else bad.
François: part 15 FCC specification - a lot of power on diode will create interference on the circuit.
Tom: sensitivity on retina of high power too (heating affect).

Preamble

Tom: Spectrix has specified 3 octets for our implementation.
Leon: is that dependent on modulation?
Tom: yes. What does it have to do: has to make some judgement about the level of the receiver signal (as gross as high or low, but maybe one of 4 categories); May have to decide to wake up some low power state things; then synchronize clock.
Richard: AC coupling?
Tom: yes, the receiver has to stabilize through the AC coupling. AGC may have to change.

Conclusion: we will just add preamble length to the modulation criteria list.

Larry: any requirement for information in the preamble other than that conveyed in the bit transitions?
Tom: personally, no.
Larry: are we talking about a PHY header that includes something besides preamble?
Tom: speed maybe.
Larry: something that identifies this as yours - whether you should wake up the MAC or not. Something beyond address - a network identifier. Or something to identify this a really data instead of noise. A Barker 11 sequence. If you don't see it you don't consider this to be anything other than crap.
David: like decoding the sync character in the PHY.
Tom: only experience may tell us whether this is worthwhile.
Larry: if there are not co-located networks and you are using this just to filter noise, network selection isn't right, don't wake up the MAC.
Maurice: remote control's do this to stop interference from other remote controls by building in manufacturers ids.
Larry: PHY level protocol information has to go somewhere - after preamble and before start of frame delimiter.
Tom: we have to decide whether we think we have as much problem as other PHYs due to our spatial limitations.

A format for these heading of modulation scheme evaluation criteria will be sent along with agenda in next mailing. Defenders of modulation schemes will bring submissions to fit those headings. Work must be done on these other issues between now and September.

Next meeting will probably be Monday night at the September interim meeting.

Meeting adjourned: 12 noon