IEEE P802.11
Wireless LANs

Tentative Minutes of meeting
Frequency Hopping PHY ad-hoc group

=======================================

Wednesday, 10 November
Frequency-hop 2.4 GHz PHY Ad Hoc Group Morning Session
Chairman Chadwick Presiding.

Peter sets the agenda as (1) finish work on the preamble, (2) review document 93/2 (Larry's draft for the PHY spec) and (3) try for at least a partial result on a unique word.

Jim McDonald: specifying an 8 microsecond ramp period allows manufacturers to make the best choices. Some implementations may have variation in when the ramp begins but good control on the slope. Others may start predictably but have slope variations.

Peter Chadwick: is it appropriate to specify this power mask on a linear or on a logarithmic basis?

Jim: favors linear

Larry Van Der Jagt: it doesn't matter. favors clarity

Jerry Socci: we could use a peak hold type instrument and look at the spectrum

Many: discuss time domain vs. frequency domain for specification and measurement. Jerry ultimately consents to time domain. Jim McDonald moves and Francois LeMaut seconds

MOTION 1: The McDonald Proposal for ramp-up (document 93/209) with language modified from dB to Watts is adopted, except the modulation during the ramp period is unspecified. For=12, Against=0, Abstentions=1. Motion 1 passed.

We discussed specifying deviation vs. specifying eyes generated by ideal receivers and decided to stick with the latter method a la DECT. Ed Geiger volunteered to be document editor for Larry's draft of the 2.4 GHz PHY section (93/172), as Francois Simon is for the whole document (92/@). We then turned to modulation during ramp up (and down). We finally decided it probably didn't matter very much. Jim McDonald moved and Francois LeMaut seconded

MOTION 2: Motion 1 is replaced with the following. The McDonald Proposal (document 93/209) modified in language from dB to Watts is adopted for both ramp up and ramp down except for the modulation pattern during the ramp period. The modulation pattern during the ramp shall be specified by the manufacturer. For=13, Against=0, Abstentions=0. Motion 2 passed.

We next discussed the training sequence, i.e., the long repetitive pattern that most cheap designs use to adjust for DC offset, find bit sync, "settle down" the receiver, etc.

Jim McDonald: wants 72 bits of 01

Ed Geiger: thinks a repeated 5 bit pattern is less likely to arise in the data. Doesn't like the 0011 pattern suggested as optimal for the "higher data rate mode" because it has only half the transitions of 01.

Larry: clock (bit sync) is generated by the PHY and supplied to the MAC.

Many: probably data will traverse the MAC/PHY interface in octets.

Ed Geiger: let's define a bit mask for jitter tolerance.

Here occurred a general discussion on bit masks.

Jim McDonald: We may want to use criteria other than power for diversity management. The 01 pattern has the shortest period and thus the fastest aquisition of register (to the repeating pattern).

Roger Jellicoe: questions Ed Geiger on the utility of bit masks.

Ed: wonders whether resynchronizing is necessary after switching antennas. Says diversity management can be performed without good bit sync.

Ed & Jim: discuss simultaneous recovery functions.

Shuzo Kato: since 1 Mb/sec is the prime data rate it is OK to use the 01 pattern.

Larry Zuckerman: recovering bit sync after differentially decoding a 0011 pattern is the same as recovering synce directly on a 01 pattern.

John: suspects that the latter option is slightly better because differentially decoding generates two transitions out of one but they are not independent (or something like that).

Tim Blaney: Jim, are you saying a diversity decision requires decoding a unique word?

Jim: no, but you might want to use more than just power. We don't necessarily have forward error correction.

Tim: interferers will often supply a repeated 01 pattern.

Roger: interferers will generate most any pattern.

Tim: so the advantage of doing any decoding at all, including simply noting the existence of transitions, is to throw out ovens and direct sequence emitters?

Jim: yes.

Jerry: my approach is like Jim's. 72 may not be enough. Wants 80 bits of 01.

BREAK

Jim McDonald moved and Francois LeMaut seconded

MOTION 3: The training sequence shall be 80 bits of 01.

Ed and Jim discussed mechanisms. Jim said his design doesn't manage diversity via RSSI measurements. It looks for transitions. It knows which frequency to monitor.

Ed: Why is 01 better for you than some other pattern?

Jim (& John): One can distinguish three levels of information that might drive a diversity decision; (1) note power (2) note transitions (3) note some pattern of transitions, i.e., some word.

Ed: we expect bit jitter on the order of 200 nanoseconds due to multipath and other causes. [notetaker's note: Was Ed implying this means it will take more than a few transitions to establish good bit sync and thus good decoding and that therefore diversity management via decoding is ill-advised?]. With Apple's FEC method, RSSI is good enough.

Tim to Ed regarding Jim and Roger: they sit on their channel and constantly monitor, exercising their diversities.

Tim to Jim: you may need a longer repeating pattern.

Ed: what would be the cost of a longer pattern?

Jerry: it takes longer to recognize a longer pattern.

Ed: You can correlate in parallel.

Jerry: that's more hardware.

Roger: we don't make decisions on correlation, just on the presence of transitions.

Jim: the 01 pattern requires less current drain.

Ed: thinks picking the antenna well will be rewarded with higher throughput.

Francois: has reservations about the correctness of assumptions he infers from the discussion.

Peter: cautions that arguments which assume an AWGN channel may be erroneous.

Roger Jellicoe moves to call the question. For=7, Against=4, Abstentions=3. The motion to call the question passed.

Vote for Motion 3: For=9, Against=2, Abstentions=3. Motion 3 passed.

Tim Blaney: says that the last big decision we made, that the FH modulation type is gmsk, was voting members only.

Chair: rules we will proceed under a 1 man=1 vote rule and the draft will be updated to show the effect of motions 1 through 3.

Jim McDonald: wants to move on to the unique word. Gear-shifting is either in the MAC layer or it is in-packet. Says every packet needs a recognizable header. Proposes an complemented unique word to signal gear shifting.

Ed: how does a default-mode-only device learn when the channel has become available once again after the period of unintelligible transmission? It needs either a length (time) word in the low rate header or a low rate end word with its own retraining sequence.

A general discussion broke out. Many want in-packet rate changes with a mandatory low-rate header. The MAC group should perhaps be asked whether they prefer a length field or an end delimiter. A retraining sequence preceding a low-rate end word is not a burden on default, low-rate devices since they won't provide it.

Roger Jellicoe: How long would a length field be, including a sufficient CRC?

John McKown (the eternal TA): Hah! We can derive it by bounding the packet length as follows. We first demand that packets are brief enough that the channel doesn't change much during a single transmission, so a good diversity decision remains valid. We assume we're indoors in a standing wave pattern at 2.4 GHz. It can be shown that, with omni antennas, the features of such a pattern have a minimum size (power peak to neighboring power trough distance) of a quarter of a wavelength. The wavelength is the speed of light divided by the frequency: lambda/4 = c/4f = (3e8 meters/sec) /(4 (2.4e9 Hz)) = .0313 meters. Dividing by the maximum T-R velocity difference we wish to serve gives a minimum time to proceed from a power crest to a power trough. Assume pedestrians and slowish robots: about 1 meter/sec. That gives 31 milliseconds, peak to trough. The channel can be viewed as static for, say, a tenth of this interval, so take 3 milliseconds as a maximum packet duration. At 1 microsecond per bit, that's on the order of 3000 bits. It takes 12 bits to express lengths like that. Then add your protection.

General discussion on whether the choice between an end word and a length word is a MAC or a PHY issue. We quit for the day.

==

Thursday Morning
2.4 GHz Frequency Hopper Ad Hoc Group, 11 November 1993, Chairman Chadwick presiding.

Peter: let's discuss questions for the mac group and the high rate group.

John: Hooks for the high rate are our business, as is the question of length vs end word.

Jim: We could just demand that whatever is done by the high rate group have no impact on the baseline radio.

Peter: unhooking is the question.

Roger: We see that there are acceptable ways to do it. Perhaps we don't need to specify exactly which one today.

Peter: but it needs doing. Why not do it?

Jim: we should avoid delay waiting on other groups.

Peter vs Roger: it's unknowable what other PHY ad hoc groups will demand of the MAC group.

Roger: thinks we can proceed just knowing it's doable.

John: presents two viewgraphs as follows, with appropriate caveats. Offers a PC program which evaluates Rt free to all.

Williard's Figure of Merit for Sync Words

R_T is the total probability of incorrectly synchronizing to an overlap word, i.e., a sequence of bits which contains some of the bits of the proper sync word and the rest random data bits, divided (normalized) by the probability that a word completely made up of data bits will appear as the sync word.

$$R_T = \sum_{m=1}^{N-1} 2^m (1-p)^{m-c(m)} p^{c(m)}$$

where

N = the number of bits in the sync word

m = the number of bits of the overlap word which are actually part of the sync word, "i.e., the number of bits which overlap, or degree of overlap"

c = "the number of bits in the overlap which, as transmitted, are opposite to or conflict with bits expected in given bit positions" --- c depends on the exact bit sequence of the sync word and is thus a function of the overlap m

p = probability of bit error

The False Sync Rate

We can define two types of false sync events.

Event type 1: Sync is declared on random noise or random data with no overlap of the true sync word. With sync word length N, this happens with probability

$$P_1 = 2^{-N}$$

Event type 2: Sync is declared partly overlapping the true sync word and partly overlapping random noise or data. This is the case considered by Williard and occurs with probability

$$P_2 = R_T 2^{-N}$$

where R_T is the figure of merit discussed in P802.11/93-143.

For each type of event, the false alarm rate is the false alarm probability times the rate of trials. The total false alarm rate is then the sum of the individual rates. The trial rate for type 1 should be the rate at which words of length N are tested by the receiver, i.e., the bit rate B. The trial rate for type 2 events is, presumably, the packet rate or B/L where L is the packet length. Therefore

$$\text{false alarms per second} = BP_1 + \frac{BP_2}{L} = \frac{B}{2^N}\left(1 + \frac{R_T}{L}\right)$$

Jim: is worried how the baseline system will know the channel isn't free during the high rate transmission. Speaks for imposing a deadline Jan 94.

Francois: all the mac proposals have a length field.

Jim: draws a header: 8 bits of ramp, then 80 bits of 01, then 16 bits of unique word = 4657octal, then 8 bits of phy signaling field, then 16 bits of packet length field, then 8 bits for protection on the length field.

Peter: Are there any objections to this header? There are none.

We adjourn the ad hoc group.