IEEE P802.11

Wireless Access Method and Physical Layer Specification

IEEE 802.11/D2 Letter Ballot Response

Comments of C. A. Rypinski

LACE, Inc.
655 Redwood Highway
Mill Valley, CA 94941
Telephone: +1 415 389 6659
Facsimile: +1 415 389 6746
e-mail: rypinski@netcom.com

Abstract:

Having cast a generic NO vote on the D2 ballot, I ask that the Committee as a whole consider my reasons for this position. Those reasons are stated as briefly as I can make it on the following pages. This type of objection does not fit the procedure which is entirely optimized for heading-by-heading iterative refinement.

Those of you who think that demonstration of successful communication between a small number of units is near proof of adequacy, should consider that radio systems are very different when the density and volume of usage becomes large.

Please consider the points on their merits. If the assertions are right, there are bound to be competing products to exploit the advantages possible. If there is not a change from metered silence and primary use of peer-to-peer, the performance cannot be adequate.

Chandos A. Rypinski

Life Fellow IEEE
Fellow Radio Club of America
IEEE Centennial Medal
Avant Garde Award -- IEEE Vehicular Technology Society

IEEE 802.11/D2 Letter Ballot Response
Comments of C. A. Rypinski

SUMMARY

I do **NOT APPROVE** the IEEE P802.11/D2 proposed draft standard to be forwarded for sponsor ballot.

The brief reason is that the Standard as written will NOT provide a competent level of accuracy, capacity and reliability primarily due to faults in the use of the radio medium.

Those who believe that large scale continuous area coverage is a real requirement for a Standardized radio LAN should carefully consider these comments. As long as the access method rests on the CCA function, successful operation in this context is not possible. This is true no matter how expertly the remainder of the Standard is executed.

If this standard is converted to products, the first buyers may be pleased. However, as increasing amounts of equipment appear in areas of high concentration, it will be discovered that excess delay and transfer failure are too frequent for reliable service. It will also be learned that the throughput of one system in the midst of other like systems is far below what could have been obtained with planning based on such a model or what one system isolated and alone would carry.

Further consideration of whether the plan rests on an adequate foundation has been halted since the DFW MAC question in November '93. Now that the consequences of that choice are apparent, the incentives for reconsidering are clear or should be. I do not expect this to happen, however those who are committing resources on this position would be well advised to estimate their exposure to the described problems and the market consequences.

Substantial changes in the access method and primary mode system model would be required to change my vote to APPROVE. The difficulties, consequences and alternatives are identified and discussed below. The cited difficulties are in the model, the lower MAC and the CCA function of the PHY, and cannot be corrected by editorial changes to a few particular paragraphs.

> The minimum action to bring the system to an acceptable level would require change to an access point originated message for the channel monitoring CCA as the criteria enabling a station to transmit. If follows that the infra-structure based access model which is now permitted becomes primary. Peer-to-peer should be the permitted service within the infrastructure or a default when

such infrastructure is not present. The minimum infrastructure is a local store-and-forward repeater with address filtering and a near ceiling height antenna.

The main difficulties with the present method are:

1. Dependence on the accuracy of the clear channel assessment function, and a gross underestimate of the consequences of interference conditions that will exist with large scale use.

2. Dependence on absence of information by coding and metering of the time duration of the clear channel state for contention minimization and priority and by enabling stations to transmit if they have NOT heard a transmission indicating that the channel is already in use.

3. Absence of the substantial radio interference and coverage benefits when all stations communicate through a common point rather than directly with each other.

4. Absence of point-managed use of channel time where the point has the information to make intelligent decisions rather than fair but random choices.

The detail consequences of these difficulties are:

5. Much more complex station logic than would otherwise be required.
"... the draft is already the most complex MAC ever defined." (P802.11-95/70)

6. With high levels of channel usage, and the threshold of the CCA function increased to a point where the channel is occasionally clear, the operation will be little improved over "Aloha" where each station transmits without regard to busy status at the start of use. Nonetheless, all of the delays shown in Figure 6-6 and 6-13 will be preserved. The protocol is short on benefit when most needed for busy channel conditions.

7. Potential system peaking of channel capacity at a moderate loading point, beyond which further increase in demand diminishes the total load carried.

8. A low level of predictability for delay and transfer success in high usage environments--much lower than what could have been achieved in the circumstances.

The eventual market consequences are:

9. Inroads from competing products which pass more traffic, cost less and preempt the channel from deferral based systems.

10. Frequent disappointment and ill will from the using public as it is discovered that the 802.11 Standard is "citizen's" band level equipment when a "cellular" quality was expected. This reaction will reflect on the competence of IEEE 802.

Further explanation for the points above are covered in the following pages.

FURTHER DESCRIPTION OF THE DIFFICULTIES

The following above problems are described in greater detail below.

The present expanded model and method will served the defined mission when the ad hoc groups operate in a low interference environment and do not generate high usage traffic. It appears to serve multiple access-points for an Extended Service Area, but incompletely defines the functions necessary to make that concept work.

1. Dependence on the accuracy of the clear channel assessment function, and a gross underestimate of the interference conditions that will exist with large scale use.

If the unobstructed path service range between stations is 300 meters, the range at which interference may be caused is 600-1200 meters. Double range is four times area. If the service range actually used is much smaller, the detectable carrier power for the CCA function remains unchanged. The number of signals that may contribute to this power level is 4 to 16 times the number within the range of the service area.

The emphasis on the "hidden" transmitter problem is a misunderstanding of what happens. While it is true that a transmitter that should be heard at a peer prior to transmitting is sometimes not heard, this is infrequent because it takes a lot less signal to move the energy detector than it does to pass messages. What happens is that there is a much larger number of non-group signals which are heard and must be filtered for relevance. Filtering by raising the threshold of the CCA is an improvement some of the time. If the threshold is made high enough, the preponderance of signals heard will be ignored and there will be a higher probability of mistakenly not deferring to signals within the group. The result will then resemble no deferral at all.

This fact is not fully revealed in simulations because of simplifying assumptions that must be made to obtain a solution as discussed in a section below.

Another type of error is caused by treatment of failed transfers. If it is assumed that such service requests disappear after failure, stability at all

levels of traffic results. In fact, the originator of a failed transfer will persistently try again. If the capacity of the system is not sufficient to remove requesters at a faster rate than they enter, the backlog of frustrated users will grow and compound the number of collisions at request time.

2. Dependence on absence of information by coding and metering of the time duration of the clear channel state for contention minimization, priority and by enabling stations to transmit if they have not heard a transmission indicating that the channel is already in use.

The potential of the Standard would be greatly increased by replacing all CCA-based transmit enable functions with hub originated enabling messages.

> Experience teaches that no high traffic radio system can be reliable unless all protocol events proceed from a positive event such as the receipt of uniquely coded messages. All protocol steps must be based on the content of the digital stream, and not on parallel analog, medium specific measurements.

Time metering of the apparent idle channel state will be frustrated by too much time appearing busy, and this is the inevitable case in a loaded system. The necessary condition for a station to transmit is that a favorable path exists between source and receiving point. This is where N rather than N! paths and a higher access point antenna become indispensable.

Just as it is important to have all stations frequency hop at the same time, it will be found desirable to have all APs and all Stations transmit at the same time.

3. There is absence of the substantial radio interference and coverage benefits when all stations communicate through a common point rather than directly with each other.

This mode might be permitted but it is not required in the present plan--and hence many of the advantages are inaccessible. The peer-to-peer assumption is contrary to the results of considerable radio system experience. Modern radio systems invariably use a base station or a repeater as the common point, because they must for effective operation. If the necessity for this topology is not recognized, the system will be incompetent in radio performance regardless of the excellence of all remaining parts. The advantages relative to the effects of interference are very large. The alternative have been tried repeatedly, and the result is "citizen's band."

A peer-to-peer system requires N! radio paths, and the shared access point requires N radio paths. An access point may have a superior antenna (on the

ceiling rather than the table) which reduces the probability of obstructed paths and the transmit power required at the station. In 802.11, the AP is defined but not required. Peer-to-peer can be permitted only when no infrastructure is available, or in time intervals allocated by infrastructure.

The use of Broadcast to solve location and route problems is not in any way diminished by the use of a repeater mode. The apparent redundant transmission from the repeater when the destination station can hear the source transmission can also be eliminated as explained in P802.11/91-95.

4. Absence of point-managed use of channel time where the point has the information to make intelligent decisions rather than fair but random choices.

After association the point knows who its members are and can filter accordingly. Stations cannot substitute accumulation of this information through monitoring with any reliability. If there were request grant for longer transfers or for virtual connections, the PCF could forecast future much of traffic load for at least tens of milliseconds. This would enable limiting of harmful effects from demand in excess of capacity.

There are many more benefits of having a channel time manager acting through the access point that all associated stations can hear. Its function can be performed over interference from unassociated stations at greater distances. Many problems are resolved because the question and the answer are at the same place physically rather than separated by a radio path.

CONSEQUENCES OF THE DIFFICULTIES

5. Much more complex station logic than would otherwise be required.

"... the draft is already the most complex MAC ever defined." From P802.11-95/70 (1.)

This project should have been undertaken with infrastructure and multiple access points, and with repeaters as the primary mode for one small group. Then it would have been possible to avoid the present complex MAC in the station. The logic for a shared store-and-forward repeater with address filtering for relevance is an easy task. The consequences of insisting upon primary use of peer-to-peer where there is no assurance that every station will be heard by every other station in a group should now be obvious. Insistence on distributed logic further extends the complexity.

6. With high levels of channel usage, and the threshold of the CCA function increased to a point where the channel is occasionally clear, the

operation will be little improved over "Aloha" where each station transmits without regard to busy status at the start of use. Nonetheless, all of the delays shown in Figure 6-6 and 6-13 will be preserved. The protocol is short on benefit when most needed for busy channel conditions.

If the channels are actually low duty cycle in usage, there is little difference between competent access methods. It is only when orderly sequential use is necessary that protocol differences become critical.

Attempts have been made to devise escapes from "busy lockout" by better methods of time weighted energy detection, by adaptive CCA thresholds and by prioritizing the amount of time deferral. These methods have increased the overhead use of channel time which might be forgivable for long transfers, but which is punishing for the half of transfers that are short.

The recovered channel time from stripping the time deferrals for priority and randomizing would give enough increase in available channel time to increase the traffic carried before refinement is required. A simple Aloha (or slotted with beacon) would have higher yield than any deferral method based on received energy in a large scale system context. It is probably nearly as good to have more channel time, than more overhead for contention minimization.

Utilization can only be made high if there is a channel access manager which knows the backlog of pending requests for service, and has the capacity to allot channel time by priority and by order-of-arrival in queue.

7. Potential system peaking of channel capacity at a moderate loading point, beyond which the total load passed diminishes.

Instability and saturation are inevitable when new requests for traffic cannot be limited to a point below the service rate. With a central channel manager, backpressure can be intelligently generated making saturation behavior predictable and orderly.

8. A low level of predictability for delay and transfer success in high usage environments--much lower than what could have been achieved in the circumstances.

A system which is not sufficiently deterministic to enable capacity and delay estimates, is in danger of concealing lower than necessary performance. As shown below, there can be serious inaccuracies in simulation results because of approximations of reality.

Approximations that diminish the accuracy of simulation results

The 802.11 Committee may be giving considerable weight to some contributed simulations. They may believe that capacities are adequate, but experiences teaches that there is never enough capacity. It is just a matter of time for the load to grow to the limit that can be carried. The simulations are quite useful for experimenting with parameter choices, but the quantitative conclusion for traffic carrying capacity may be very optimistic in the context of large scale systems.

The following approximations are often made in first system plans and following simulations, and result in quantitative errors in the conclusions:

1) It is assumed that signal level above a threshold assures a very small number of failed transfers. However, because of multipath propagation effects, many signals are unreadable independently of level. Though a minimum signal level is required for predominantly error free transmission, this is not a sufficient condition. There is a significant probability of flawed transfers at any signal level when multipath effects are fully considered. This problem is exacerbated by N! necessary propagation paths rather than N particularly when one end of the N paths can be a superior antenna.

2) It is assumed that the Rayleigh distribution is accurate across the full range. This function is an "approximate but good" fit to values between 5% and 95% probabilities. Improbable levels are a larger fraction of all levels received than is given by the Rayleigh function.

A small proportion of messages (a few %) will come from unexpectedly distant points in part due to geometry's in the environment causing reflector focusing or waveguide type propagation. What is a small proportion in a signal level population is a large proportion for failed transfers.

3) A third inaccuracy that occurs is in summing the effect of multiple interfering signals. On the average a power sum is correct, but the ability to cause errors may be more dependent on instantaneous voltage sum. Some errors will be caused by peak values. The power sum will underestimate the error causing power.

Incidentally, the addition of multiple statistically described signals into a composite function seems to be a very difficult problem. This would be very helpful for describing interference in signal-to-interference ratio.

CONCLUSION

Failure at any time to reconsider the original DFW MAC decision, has resulted in a document with flaws that will limit its useful lifetime at the least. This is unfortunate since a large portion of the document is expertly produced and would be reusable with an appropriate radio MAC. Instead, it is likely that it will be bring disappointment in the institution and individuals who produced it.

From a procedure view point, there should have been an opportunity for overall review sometime this year. I would press that point, except that I think it futile. I do not have the time and energy for the political and technical effort required to put forward the alternatives in document and test result form which might win such a constituency. That is too bad for future purchasers of 802.11 based products.

If these changes are to be made, they will have to be made by others in this group. In the unlikely event that such a decision is made, I would help.

Chandos A. Rypinski