One Approach to Wireless Network Architecture.

Document 802.11/91-2

Dave Bagby

4 January 1991

Introduction:

This paper describes one approach to the definition of a wireless network architecture.

This paper is to be distributed at the January 1991, 802.11 meeting, where I hope it will act as fuel for discussion of the functional requirements of a wireless network.

The paper is written in an informal style. (I apologize for the rough edges - much of this was written at home where I don't have access to all the fancy document tools.) Within this paper I will:

- Examine some of the unstated assumptions heard during previous meetings.
- Define some important terms and their relationships.
- Describe an approach for defining a wireless LAN system.

Philosophy:

During the first two 802.11 meetings, there were many (sometimes heated) discussions of frequency bands, propagation characteristics and modulation techniques. Various members of the committee appear to have a favorite radio in mind during these discussions. It seems that the assumption of a particular radio often drives the discussion of a system architecture.

This is the "bottom up" approach to design. Too much emphasis is being placed on the physical layer at the expense of overall system architecture.

I may be in a minority; I am primarily a computer engineer. My expertise is in computer hardware and software design. My radio experience is primarily that of a radio user (rather than a RF engineer). I don't think that the physical layer issues are most important at this stage of our work.

I recommend we take a top down approach to our task;

First, define the functionality we require of a wireless network,

Second, examine the available wireless technologies, Third, pick the technology which best satisfies our functional requirements.

Motivation by previous work:

Portions of my approach to defining a system architecture were prompted by Chandos Rypinski's paper (doc # 90-18) "Wireless System Architecture - Major Choices and Considerations". That paper examines some issues within the physical layer and derives an architecture based on the tradeoffs made. (My comments on document 90-18 can be found in appendix A of this paper.)

While reading document 90-18, I began to consider the functionality a wireless system should provide, irrespective of the physical medium chosen. I decided I needed to remove myself from the details of any specific PHY layer and consider the larger picture.

Assumptions and the 802.11 PAR:

During the previous two 802.11 meetings we have managed to generate one document, representing an initial consensus about what we are trying to accomplish. (90-19: The proposal for an improved PAR, as submitted to the executive committee in November.)

I assume (in the absence of additional information) that the PAR was approved by the executive committee as submitted. If significant changes were made by the executive committee, they will require discussion during our January meeting.

It seemed reasonable that any architectural approach presented MUST satisfy the requirements contained in the PAR. I began by examining the constraints contained within the 802.11 PAR.

The PAR contains several subsections, each of which has something to say about the requirements of a wireless network.

Type of medium:
The section specifies the use of electromagnetic waves through the air, and says we will define at least one PHY layer within the standard.

The implication is that the MAC must have the ability to handle possible multiple physical layers. I believe this to mean "multiple alternative physical layers", i.e. the MAC only handles one at a time. Another interpretation is that the MAC must handle multiple physical layers simultaneously. This issue should be addressed by the committee.

The ideas presented in this paper are not particularly dependent upon the frequency of the electromagnetic waves chosen for the PHY layer. I favor a PHY layer based on RF technology and have made that assumption within this paper.

The issue of RF vs non-RF should be addressed by the committee during discussions of alternative PHYs.

Radio Spectrum:
The section states that our initial effort will be for the ISM bands and that we will consider the use of additional bands.

This implies it should be possible to implement the system architecture with a PHY layer based on the ISM bands. This is true for the proposals contained in this paper.

Supported Stations:
The section lists several categories of example stations which are to be supported.

Nothing in this paper would prevent any of the listed station types from being supported.

Environment:
The section refers to several concepts which are not defined. The concepts are:

1) Stations.
2) Basic Service Area (BSA).
3) Distribution System.
4) Extended Service Area (ESA).

This paper will offer definitions for these concepts and derive some simple relationships between them. A sample distributed architecture using these definitions is then described.

Supported Service:
The section requires a connectionless service and a packetized voice service.

In my opinion, the primary use for a wireless network system is to interconnect a nonhomogeneous set of computers without the use of wires. This task is achievable, in need of standardization and of significant commercial value.

The requirements implied by real time voice communication between people, complicate the system out of proportion to the utility gained. Such a system could be supported, but the complexity needed for reasonable performance appears high.

I am assuming that the requirement for packetized voice does not mean real time, interactive, voice communication. Instead, voice service is assumed to refer to non-real time voice messages; something similar to a voice mail box.

This is a fundamental functional requirements issue, and should be addressed by the committee as soon as possible.

Compatibility requirements:

This section calls out several other standards which our work shall be compatible with. I believe that nothing in this paper is in conflict with the referenced documents.

Additional Assumptions:

Additional assumptions made which have flavored the paper are:

The system must be robust in design.
Single point failures which take down an entire network are not acceptable. Error conditions should result in performance degradation rather then network failure whenever possible. This leads me to consider a decentralized network design.

The system design should be driven by market requirements.
We should bias our trade offs in favor of the high volume markets. (While doing everything possible to accommodate generality and niche applications.) Technical elegance should not be achieved at the expense of commercial viability. I do not wish to invest the time necessary to create a new standard, only to see it be a commercial flop.

Ease of installation is important.
Wireless networks will not necessarily (and probably won't) be installed by traditional MIS departments. A typical PC user should be able (and willing to attempt) to set up two computers (using a product based on the 802.11 standard,) and make use of a wireless network. This must be true for BSA networks and preferably true for ESA networks.

Ease of use is important.
A wireless network should be no more visible to a station user than a wired network. Ideally a user would not be able to tell whether he is using a wired or wireless MAC.

Hardware costs are important.
Consider the cost trends in station hardware. Wireless network hardware should be some small percentage of the cost of the station it is going into.

Operating costs must also be considered. The operating cost I am most concerned about is power consumption. I consider it vital that battery operated station adaptors be available. Disallowing battery operation would remove portability, defeating a major purpose of a wireless network, eliminating many types of mobile stations.

I may be assuming something of which I am unaware. If so, please let me know, that's how I learn. (Besides, I don't worry that this committee will let any significant assumption go unexamined...)

Some Working Definitions:

Earlier in this paper I noted four concepts from the PAR which required definition:

1) Stations.
2) Basic Service Area (BSA).
3) Distribution System.
4) Extended Service Area (ESA).

The station:

I will start with a working definition for a station.

Definition 1:
A station is any computer which contains an implementation of an 802.11 MAC and PHY.

The packaging of the 802.11 implementation is not important. For visualization, let's assume the implementation is a card or module which is inserted into the station computer hardware (it could just as easily be built in permanently, it really doesn't matter). Let's call this module the station adaptor.

The computer uses the station adaptor to exchange information with other computers with station adaptors. The station definition is intentionally broad, any computer with a station adaptor can be a station.

Note that there are no requirements for a person to be at a station. This is important. The architecture should not assume human interaction at a station. This will be an important principle when we consider error recovery issues.

The Basic Service Area (BSA):

The PAR says that a BSA is that area "... in which each station can communicate with any other station in the BSA".

I interpret the requirement for a BSA to be that area where each station can talk *directly* to another station. If station communications within a BSA were not direct, there would not be a distinction in the PAR between BSA and ESA. Nor would there be a need for the concept of a Distribution System (which the PAR states exists to handle physical ranges larger that of a BSA).

Definition 2:
A Basic Service Area (BSA) is that area in which each station can directly communicate with any other station within the BSA.

I assume the use of omnidirectional antennas (at least in the horizontal plane). I further assume that all stations have the same range, and that a station's transmit and receive ranges are equal.

While equal range among stations is not strictly necessary, I see no advantage to allowing different stations (with standardized PHY layers) to have differing (nominal) ranges, nor do I see a benefit in staggering the receive / transmit ranges.

Omnidirectional station adaptor antennas generate circular coverage areas centered about the station. I call the radius of this circle Rs (for station range).

Some simple geometry will show that the worst case for two stations in communication, is when the stations are on opposite points of the circle's circumference. Since each station must be able to reach the other directly, the limit to the diameter of the covered area is Rs. See figure 1.

Figure 1

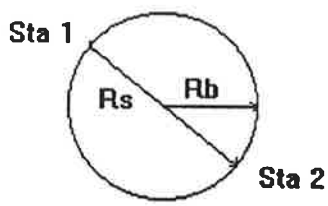

BSA Geometry

Formula 1:
Given

Rs = the radius of the area covered by the station
Rb = the radius of a BSA

*Rs = 2*Rb*

Formula 2:
Given a station range of Rs, the area of a Basic Service Area is given by:

*Ab = Pi*Rb^2 = Pi*(Rs/2)^2*

This gives us the first simple relationships between station range and BSA.

Once a minimum BSA size is decided, we will know the range required of our PHY layer.

(For sake of discussion, I have concentrated on the 2D cases to illustrate the principles involved. My principal interest is in the office environment market, which is adequately described by the 2D approximation as long as Rs is sufficiently larger than the office ceiling-to-floor distance. This should hold for real life, useful values of Rs.

Those people interested in factory environments should consider values of Rs which account for high ceilings. they should also think in terms of truncated spheres instead of circles.

Perceived market requirements will cause one to favor one value of Rs over another (and hence the choice of PHY technology). This is one reason I am NOT discussing magnitudes within this paper - we need to understand the relationships before getting mired in a discussion of specific values.)

The Extended Service Area (ESA):

The PAR states that "Stations which interoperate in both BSA and ESA shall be defined if feasible".

It is feasible to define such a station and I have done so in this document. I consider the ability to use identical station hardware in both a BSA and ESA network crucial to the success of the standard.

The PAR indicates that a "Distribution System" is to be defined which is "... designed to provide range extensibility...". There are several approaches to a distribution system. The discussions I've heard have all been variations of the cellular concepts.

In a cellular type system, the distribution system provides cells of coverage area. A station operates within one or more of these cells.

I will take the liberty of redefining a term from Chandros' paper (90-18):

Definition 3:
Access point is the term used for a fixed radiation point provided by a Distribution System.

The changes from Chandros' definition are small but significant.

The phrase "... a fixed radiation point..." is used rather than "... the fixed radiation point..." because I want to make it explicit that a distribution system can provide one or more access points.

I used "distribution system" rather than "... wired system..." because I want the definition of Access Point to be independent of any particular type of distribution system.

Now we can define an ESA:

Definition 4:

An Extended Service Area (ESA) is the sum of the coverage areas provided by a Distribution System's Access Points.

I again assume omnidirectional antennas, implying a circular cell shape centered around an access point, and that the range of an access point be equal to the range of a station adaptor.

This yields another relationship:

Formula 3:
Let Ra = Range of the access point, then
Ra = Rs

While this is not strictly required, it has the following advantages:

1) It is simple and meets our needs.
2) It allows station hardware to easily operate in both a BSA and ESA environment.
3) It allows portions of the station hardware and the access point hardware to be identical. This reduces complexity and increases manufacturing economies of scale, which in turn helps reduce the cost of station adaptors.

The Distribution System:

I offer the following definition for a Distribution System.

Definition 5:
A Distribution System is defined to be that system, which links a set of Access Points together, in such a way that stations within different Access Point coverage areas, can communicate.

Note that the definitions do not have anything to say about whether cell coverages are, or are not, contiguous. There will be installations where contiguous coverage is not desired. There will also be installations where contiguous coverage is required.

To provide contiguous coverage over an arbitrary area, we must tile the area to be covered with access points. The two obvious candidates for tile shapes are squares and hexagons. This paper was written using square tiles for the analysis.

Unfortunately, we have circular coverage areas (anyone know how to make a square pattern antenna?). The most efficient tile size will be that tile, which is inscribed by our circular coverage area. See figure 2.

(I believe the square to be sufficient for the purposes of this paper. In the future I may redo the formulas based on hexagons. Later in this paper I describe two possible distribution systems; one fits square tiles nicely while the other may operate more efficiently with hex tiling.)

Note that any tiling system will have areas which are covered by more than one access point. This is unavoidable and the protocols used within the Distribution System have to be able to handle this.

With circular cells areas, and square tiles, we can achieve contiguous coverage by spacing access points such that the diagonal of the tile is equal to the access point coverage radius. See figure 2.

Figure 2

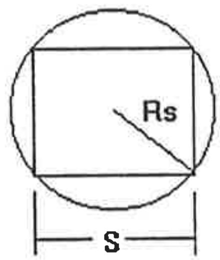

ESA Cell Geometry

Since we know that Ra = Rs, we can derive the size of the tile and hence the required spacing of access points for contiguous coverage.

Formula 4:
Let the length of the side of the square tile be given by S.
Then

    2*(S^2) = (2*Rs)^2

or

    *S = Sqrt(2)*Rs*

Because the tiles are squares, S is not only the length of the side of the tile, it is also the spacing required of access points for contiguous area coverage.

The area of an access point's exclusive coverage is given by:

Formula 5:
Let Ac = the area covered by an access point (exclusive of cell overlap).

    Ac = [Sqrt(2)*Rs]^2

    *Ac = 2*Rs^2*

Now we can derive a relationship between ESA cell area (Ac) and BSA area (Ab).

From formula 5:

    Ac = 2*Rs^2

giving

    Ac/2 = Rs^2

```
From Formula 2:
     Ab = pi*(Rs/2)^2
giving
     4*Ab/pi = Rs^2

therefore
     Ac/2 = 4*Ab/pi
and
```

Formula 6:
```
     Ac = (8/pi)*Ab
```

Using these basic relationships, I offer the following tables as food for thought. Table 1 fixes Rs and gives the resulting BSA area and ESA cell area. Table 2 fixes a BSA size and gives the Rs required and the corresponding ESA cell size.

Table 1

Fix Rs:

Rs (ft)	Ab (sq ft)	Ac (sq ft)
25	491	1250
50	1963	5000
100	7854	20000
200	31416	80000
300	70686	180000
500	196350	500000
1000	785398	2000000

Table 2

Fix Ab:

Ab (sq ft)	Rs (ft)	Ac (sq ft)
500	25	1273
1000	36	2546
2500	56	6366
5000	80	12732
10000	113	25465
25000	178	63662
50000	252	127324

I found it very instructive to spend some time looking at these tables and looking at the sizes and layouts of various office buildings. I highly recommend that members of the committee spend some time considering their customer's environments, visualizing access point based on several values of Rs.

Consider the relationship between BSA and ESA sizes and contemplate the number of access points needed for a given building layout. Practical limits will probably be determined by the dimensions of unobstructed areas within the building (rather than ESA cell sizes). I suspect that no one wants to try and sell a system which requires access points every couple of feet.

As usual the various desires are in conflict and must be balanced. As the subject of a separate paper I intend to address the factors which lead to my personal choice for a value of Rs.

I point out that there is no analysis here of data bandwidth. Bandwidth issues arise during discussion of PHY layers. I assume everyone wants the highest bandwidth possible, consistent with the limitations imposed by complexity, power consumption and the

FCC. I believe that bandwidth issues can be treated as orthogonal to the issues raised within this paper.

Some Distribution system approaches:

During the remainder of this paper I am less quantified about the information presented. My purpose is to provide enough information to allow the reader to grasp the concepts and begin to think about the pros and cons. I make no attempt to give the details necessary for either complete specification or implementation of the concepts.

Next I describe a distribution system architecture. The system is distributed in nature, with the system intelligence contained in the access points.

A distribution system must link together one or more access points. The method used to link the access points must provide communication between those access points. The communication is done over some type of media. I call this the Distribution system Media.

Definition 6:
The Distribution System Media (DSM) is the media used by a Distribution System to interconnect access points.

The important concept is that there are two different *logical* media which we are using for communications; the 802.11 PHY media (PM) and the distribution system media (DSM).

The physical DSM will probably be different than the PM. However, the logical model intentionally does not require this.

An additional requirement on the distribution system (from the PAR) is that it be "managed". Managed is not defined by the PAR. In a practical sense, "managed" means to be able to control the behavior of the distribution system.

Control is accomplished by providing input to the program which runs the distribution system algorithms. In a distributed system, that intelligence is contained within the system's access points. To manage the access points, they must be addressable. The proposed architecture provides for this.

Let's consider an access point, what is it?

An access point must interface to the PHY layer media, and it must interface to the distribution system media. The access point's primary job is to translate messages between medias. It seems reasonable to assume some intelligence to be necessary for this translation (the intelligence needed is addressed later).

Figure 3 contains a functional block diagram of an access point.

Figure 3

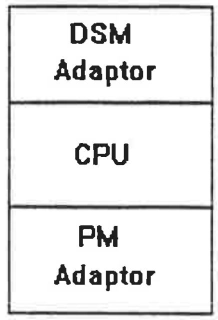

Access Point Functional Blocks

With this functionality in the access point, let's investigate a distribution system architecture. One significant distribution system variable is the choice of the DSM. The DSM forms the spine of the distribution system.

Wired DSM:

There will be many cases where a wireless LAN will coexist with a wired LAN. In these cases it would be economical to make use of existing building wiring for the DSM.

I refer to this situation as a wired DSM. A wired DSM could result in the example distribution system in figure 4.

Figure 4

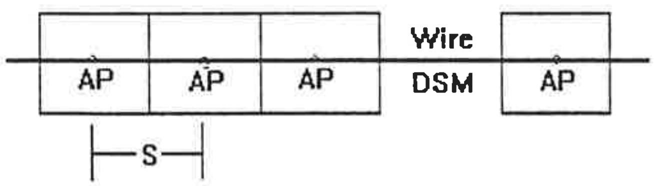

An ESA network using a Wired DSM

This distribution system has two distinct coverage areas. In the first coverage area three access points are spaced along the DSM cable with spacing equal to S, providing a contiguous area of coverage (of total area 3*S^2 and dimensions S by 3S). A single cell of additional coverage is also connected to the DSM cable, providing a logical extension of the 802.11 LAN into a separate physical location.

This distribution system has several interesting features:

Failure modes:
The failure of an access point will only affect those stations within the coverage area of the failed access point - they become isolated from the rest of the network. It is possible to design the network software so that communications between stations

within a single ESA cell are direct (rather than through the access point). Failure of the access point hardware will only be noticed by those stations attempting communication via the failed access point to other cells.

Figure 4 assumes the use of the minimal number of access points necessary for contiguous coverage. If an extremely robust system is desired, the spacing between access points could be less than S. For example, at a spacing of S/2 (ignoring cells at the edges of an ESA tile pattern) each station would always be covered by two access points and the failure of any single access point would not be noticed.

Note that distribution system operation is independent of access point spacing. The spacing may be adjusted by the installer to provide the desired redundancy. There are performance implications related to extreme cell overlap. The lack of sensitivity to access point spacing also contributes to ease of installation since each access point does not have to be precisely located.

Failures of the DSM cable only serve to segment the LAN into smaller LANs, each of which are still operational. Distribution system algorithms can be devised so that repair of the cable will simply rejoin the severed sub-nets (dynamically, without network reinitialization).

Performance issues:
In many cell based radio systems multiple frequencies are used to separate adjacent cells. Both four frequency and seven frequency patterns are used (depending the amount of cell isolation desired).

With this approach, station adaptors must have multiple logical channels. Multiple channels are achieved at the expense of data bandwidth.

An alternative, (which I currently favor,) is to use the same channel for all cells, allowing the radio module to be single channel. (One may want more than one physical channel in the radio to implement multiple physical networks, but that's a separate discussion). The use of a single logical channel simplifies installation and configuration of access points. With a single channel approach, access points act as time domain repeaters between PM and DSM.

An access point is dealing with two different medias (PM and DSM). Therefore, it seems likely that there is some message buffering within the access point. The use of the time domain appears to be a natural fit.

Traffic volumes:
One might be tempted to assume that all messages which are received by one access point, are sent via the DSM to all other

access points for retransmission. This would result in unacceptable amounts of overhead traffic within the network.

Access points must contain intelligent routing abilities. The routing ability is hidden within the distribution system and is not available for use outside the system. An access point can look at a message received from a station within it's coverage area, and forward the message to the access point appropriate to the intended destination station.

There are established algorithms which handle this type of routing situation. The same dynamic algorithms will automatically handle station messages received by multiple access points, as well as mobile stations moving from one cell to another.

An architectural advantage:
Consider the access point functional block diagram given in figure 3. For all the functions discussed, the physical nature of the DSM is irrelevant.

The interface, from the access point processor to the DSM simply has to be able to deliver a message from one access point to another. The basic requirement is that access points to be addressable entities on the DSM. Several existing wired networks satisfy this requirement.

In fact, an interface from the access point processor to a DSM could be defined which is independent of the physical DSM used.

This is a significant advantage. It eliminates the need for 802.11 to specify a media for the DSM. It provides opportunity for product differentiation by allowing companies to provide different performance level distribution systems while remaining 802.11 compliant. If defined correctly, it should be possible for station adaptors from vendor A, to talk to vendors B's adaptors and access points (and therefore B's distribution system).

Cost issues:
It may seem that each access point is asked to do so much that it will be an expensive piece of hardware. I don't believe this to be the case. Access points will cost more than station adaptors. For an ESA network, there are few access points and many station adaptors. Overall system cost should be quite reasonable.

This is an area where the choice of Rs affects system cost. If you need an access point every few feet, the system cost gets high. Via an appropriate choice of Rs, the number of access point modules can be kept low (even for large facilities).

Other DSM choices:

There is no reason that a DSM must be wire based. A wireless DSM would enable the installation of ESA networks in facilities where wires can not be run. I will not explore this area in more detail here other than to note that a wireless DSM might be better modeled using hexagonal tiles. I will be happy to discuss the implications of a wireless DSM with interested parties.

Summary:

This paper has

- Provided working definitions for portions of the 802.11 architecture.

- Identified several issues which the committee needs to address.

- Presented a general architectural approach which
    - a) Is distributed in nature.
    - b) Provides a robust system design.
    - c) Is PHY layer independent.
    - d) Considers all requirements contained within the 802.11 PAR.

I look forward to discussing these concepts during our January 1991 meeting (Step right up folks, 3 shots for a dollar...)

This page has been left nearly blank intentionally.

Appendix A

Comments on Document 802.11/90-18

Document 90-18 contains some excellent discussion of PHY layer design issues related to station range. The paper argues for a design which minimizes reach and recommends a centralized control architecture.

The paper starts out by stating several premises. As stated in the paper, some of the premises are arguable.

I wish to address several points contained in Document 90-18.

A word about format; Quoted paragraphs were extracted from document 90-18. My comments pertaining to the issue immediately follow the quote.

LAN capacity and speed:

"LAN capacity and speed capability approaching or equaling that of existing wired LANs."

I agree with the sentiment, but disagree with making this a requirement. The specific target (of Wired LAN performance) may or may not be appropriate. Without functional requirements and target market definitions, it is impossible to judge.

Given the complex issues concerning spectrum space, modulation techniques and bandwidth, it would be a mistake to fix the bandwidth at a fixed value this soon.

Use of EIA/TIA 568 wiring:

"Dependence upon and compatibility with the telephone twisted pair wiring defined in EIA/TIA 568 for operation within user controlled premises."

There are significant problems with this premise:

Very few buildings are wired to EIA/TIA 568. Making the 802.11 standard dependent on this specific wiring scheme significantly jeopardizes market acceptance of the standard.

There is no reason for the 802.11 standard to be dependent on any particular building wiring. In fact, it seems odd to require a wired infrastructure for a wireless LAN.

The referenced wiring standard would mandate a physical topology based on wiring closets layouts. This may not match requirements for wireless LAN coverage areas and is too limiting.

Acquisition of spectrum and the use of minimal reach:

"Non-exclusive co-use of radio frequencies now occupied by point-to-point microwave radio links."

This premise drives the architectural contents of document 90-18.

I understand the strong desire to obtain protected spectrum space for wireless LAN implementation.

On page 3, 90-18 states in bold letters: "Above all: Unless reach is minimized, the potential for interference to existing 1.7 - 2.3 GHz services may prevent FCC authorization of these frequencies for many years".

The converse of the quoted statement is NOT true. Minimizing reach will NOT cause authorization of use of the stated frequencies.

The 802.11 PAR requires that our *initial* efforts be for ISM frequencies.

The PAR also authorizes the committee to lobby for new frequency allocations. The efforts to acquire frequency allocation are fine.

Allowing a design goal (minimum reach), which is primarily motivated by the desire to acquire specific spectrum space, to be the driving factor which defines the system architecture, is not a good approach. It makes the system definition dependent upon a specific spectral resource which is not available. That is not acceptable.

The PAR requires that the MAC be able to accommodate multiple PHYs. Not all PHYs will have minimum reach as a primary design goal.

The architecture described in 90-18 is too heavily influenced by the desire to have an architectural carrot to hold in front of the FCC while lobbying for spectrum allocation.

Because of the use of minimal reach, the described architecture requires the use of a distribution system to achieve any reasonable coverage area. The paper appears to ignore the PAR requirement that the standard address both BSA and ESA networks. At best, it seems that the BSA area size would be so small as to be useless (See relationship between BSA size and station range in document 91-2).

Optical Propagation:

The paper states that the shorter the reach, the greater the probability of optical path Propagation. I heard this questioned during the November 802.11 meeting. The assumption of optical path Propagation appears to be key to portions of the paper's analysis. The issue should be addressed further by the committee.

Centralized Control:

The paper promotes a network based on a central controller.

"... mandatory for access-point receivers to report signal level to a central controller to enable implementations of smarter algorithms..." (page 5)

The stated desire is for the system to be able to make use of signal level information available at access points. The entity which process information is independent of the source of the information. Centralized control is not required.

"The grant of access for stations must be centrally controlled contrary to the philosophy of other 802 LANs. The alternative is to broadcast all transmissions to all stations and to provide in each the logic to determine that its transmission is permitted." (page 6)

The embedded implication is that the only alternative to centralized control causes station hardware to increase in complexity and become prohibitively expensive.

Fortunately the broadcasting of all transmissions to all stations is NOT the only alternative. (See document 91-2 for additional alternatives). Centralized grant of access is not required, and this is not a justification for centralized controller.

"Minimum logic and cost in stations is obtained by locating access control centrally." (page 6)

This is not true. Station cost is a function of more than the cost of station adaptor logic. In some cases, removing logic from the adaptor can actually increase station costs (due to processing burden on station).

Station *adaptor* logic can be minimized by removing logic from the station adaptor. There is no reason the shifted logic must reside in a central network controller.

"Stations transmit only after a 'permission' message from the central controller. Stations know nothing about the overall status of the system." (page 6)

This is a polling based network access scheme. I have concerns over performance issues associated with polling schemes. The station user must not be aware of limitation on when the station can transmit. This requires that polling must be fast enough to be invisible. It may be very difficult, if not impossible, to achieve this with large numbers of stations.

It is implied that it is undesirable for a station to have any knowledge of system status. This is used as justification for permission messages. I don't understand why station knowledge of system status is inherently bad.

Under the section "Message-based access control" (page 6), several functions are listed which are proposed for the central controller. None of the listed functions are endemic to a centralized controller.

Two of the listed functions raise additional architectural concern.

"Operate a background poll to detect deactivated stations which have not informed the controller of that change by message."

From the station viewpoint this imposes what I call an active idle state. It requires the station to respond to periodic health and welfare inquires, (even if the station was not using the network at the time the inquiry was received). The act of answering the inquiry requires the station to transmit and it is transmitting which consumes the majority of station adaptor power.

This is wasteful and a burden to battery operated stations.

Further, active idle states are not necessary. There is no need for the network to detect inactive stations. Active idle protocols place undesired restrictions on station operation and adaptor implementation.

Consider the user of a PC which contains a wireless network station adaptor. The following sequence is very common:

a) Access the network to get some needed information.
b) Perform local work which does not involve network access.
c) Access the network again.

The time period between a) and c) can often be hours long. Why should the station user, station adaptor or network controller care about the time in the middle? There is no error condition until the network tries to access a deactivated station - at which time you have the same situation as if the station had failed to respond to a health and welfare inquiry.

The majority of PCs run a single tasking operating system. Low cost implementations of station adaptors are likely to use the host processor to manage the network adaptor. The use of active idle protocols makes this more difficult than it need be.

"to solicit requests for access from stations and to grant permission for stations to send datagrams..."

This appears to be a multiple step process:

a) Ask a station if it wants to transmit
b) The station says yes
c) Tell the station it can transmit
d) The station transmits.
e) The next station is polled.

Assuming a polling access scheme must be used, it would be more economical to:

a) Inform a station that it can transmit if it wants
b) The station sends either data or a "no thank you"
c) The next station is polled.

Document 90-18 does not attempt to address issues associated with failure modes. I point out that centralized controllers inherently create the possibility for a single point failure to disable the entire network. The markets I deal with will not accept this.

Conclusions:

- Counter arguments exist for several key premises.

- The architecture presented is overly influenced by FCC political considerations.

- No functionality endemic to centralized control was offered.

- The use of centralized control contains significant disadvantages.

Centralized control of a network does not allow an architecture with sufficient functionality for wide spread market acceptance.

The architecture described in document 90-18 can be improved upon and should not be adopted as the 802.11 architectural model.