A UNIFYING APPROACH: CENTRALLY MANAGED VS PEER-TO-PEER TOPOLOGIES

Juan Grau
295 N. Bernardo Av.
Mt. View, CA 94043

P r o x i m

OVERVIEW

Arguments have been made in support of both centrally managed and peer-to-peer or distributed media access methods. Both have advantages under certain conditions and suffer from limitations in other circumstances.

In deciding on a MAC layer standard that will satisfy the broad range of present and future applications possible with wireless LAN technology, this committee should consider the impact that such a standard will have in addressing these applications. In addition, the approved standard should allow room for growth, as undoubtedly we will not be able to foresee all the needs of the future. This objective, however, must be balanced against the practical aspects of implementing wireless LANs today. Real constraints imposed by regulatory environments, technological limits and even political realities among members of the committee all need to be carefully considered. To date, this committee has officially embraced one PHY layer standard: a frequency hopping PHY layer. The eventual MAC layer standard must address performance, cost, ease of use and other factors which will affect customer acceptance in the context of the approved frequency hopping PHY layer.

It is important to remember that mobility has been one of the strongest forces driving the move towards wireless products and services. In drafting a standard, this committee should have a clear understanding of how this standard will address the myriad of MAC layer issues associated with mobility such as power management, roaming, intermittent connectivity to the network, and ad-hoc networks, among others. There has been a lot of theoretical work done in support of the different proposals, unfortunately not all of this work has been substantiated with empirical real world data. Only a handful of vendors have had systems in operation long enough to be able to put their hypotheses to test. Moreover, the few systems that have found their way into the marketplace have generally been deployed for different applications, making it more challenging to find a common denominator among them. It is important, however, that this common denominator be well defined if there is any hope for a successful standard.

In the following paragraphs, we will highlight unifying features of the different proposals now on the table that we deem to be key to the eventual standard. Our comments are based on several years of market-based experience and the knowledge gained during the development of two commercial spread spectrum systems: one using direct sequence and the other frequency hopping technology, one implementing a pure peer-to-peer topology (RTS-CTS-DATA-ACK with reservation type of protocol), the other relying on selected elements of a centrally managed architecture.

CENTRALLY MANAGED VS. DISTRIBUTED NETWORK TOPOLOGIES:

The following table summarizes some of the key attributes traditionally associated with each alternative. It is followed with a brief discussion of their potential impact on practical systems.

Centrally Managed	Distributed
Easier synchronization Easier power management Easier "other" management Good performance for isochronous traffic Predictable range Predictable latency Predictable performance under heavy load	Ease of configuration Ease of installation Ad-hoc support Support for collocated independent nets Good performance for bursty traffic Simpler MAC -> Less expensive

It should be noted that it is possible for a unified MAC layer architecture to address the limitations traditionally associated with each topology. This is in fact the direction that several of the proposals now on the table have taken. However, the compromises that this leads to must be evaluated thoroughly against the objectives targeted, performance and otherwise. This is especially true in the case were the proposals are not backed by fully deployed systems.

Ease of synchronization:

This category extends beyond the management of frequency hopping timing addressed in previous proposals, although synchronization is clearly critical to timing in frequency hopping environments and is more easily accomplished with elements of a centrally managed approach. Some activities that require synchronization between members of a wireless network include:

1) Power management: The remote station will periodically put itself to sleep to conserve power. This periodicity requires that a time reference be kept and preferably managed by a central agent.
2) Isochronous Traffic: Periodicity of transmissions require that a time reference be kept
3) Frequency Hopping: Hopping information needs to be available to all members of the network that wish to communicate with each other.

Although it is possible to devise ways to synchronize on a distributed basis, this is a much more difficult proposition and will result in significant performance degradation whenever synchronization is lost. This will undoubtedly be the case in mobile networks as stations roam over an extended service area. An additional difficulty exists with the possibility of split networks. This is a situation in which a subset of the members of a network are synchronized to a different reference from the rest of the members of the network and cannot communicate with them.

When a station roams from one basic service area to another, it will have search for a new time reference and synchronize to it. This "hand-off" will need to appear seamless to the user. Consequently the re-synchronization function should be designed accordingly. Performance of an eventual standard under these conditions needs to be clearly understood.

Central management is a more effective way to handle all synchronization functions. The eventual standard should be evaluated on its ability to handle synchronization functions effectively.

Ease of Installation/Configuration

The plug-and-play potential of peer-to-peer topologies brings a level of user friendliness that is hard to match with a centrally managed system. By nature, centrally managed protocols require the user to have some understanding of the network topology. This is quite acceptable for medium to large size networks, where a knowledgeable network administrator will likely be available. It is not optimal for very small networks were non-technical people will likely have to install and manage the network (e.g. the doctor's office, the travel agency, etc.). A significant portion of the growth in the wireless networking business is likely to come from extending networking services to consumers that do not currently use networks. The easier wireless LANs are to install and use, the greater the market penetration of this critical technology.

Peer-to-peer networks have traditionally been easier to install and configure. The eventual standard should accommodate implementations requiring little or no knowledge of the network topology, including support for Ad-hoc networks.

Co-located independent networks

Requiring *a priori* knowledge of the network topology poses a problem when independent networks (i.e. networks belonging to separate entities) share a common geography. This situation may arise, for example, when the two entities share halves of a building, separated perhaps only by an aisle or a floor, or when an ad-hoc group, such as an external audit team, is established inside the premises of an existing wireless LAN installation. Peer-to-peer topologies can handle these situations more naturally than centrally managed ones. It would be possible to allow an existing central management agent to provide the important management feature for everyone in its coverage area. However, it would be difficult to persuade a customer to allow its network to be managed by a base station belonging to the company next door.

The existence of co-located independent networks is a reality that we cannot ignore. The eventual standard needs to clearly identify how it will deal with such situations.

A FLEXIBLE, UNIFYING APPROACH

There are two proposals currently under consideration by this committee that incorporate many of the elements needed for a workable standard: the R-TDMA proposal and the WHAT proposal. The first one approaches the challenge of defining a MAC layer from a perspective that a centrally managed topology is better, the second one takes the opposite view. Neither one addresses in detail all the issues discussed above. It is possible, however, to bridge the gap between these two proposals in a flexible way, minimizing the level of compromises required to achieve true interoperability.

The eventual IEEE standard should combine elements of both centralized and distributed proposals in a flexible way:

1) Selected central management functions. The MAC layer standard should require centrally managed timing reference at a minimum. It should also carry and distribute information about the configuration of the protocol. In the absence of any configuration information, the protocol would revert to a pure RTS-CTS-DATA-ACK type of protocol that has already been proven in the market. Minimizing the required management functions to just a simple timing reference will alleviate problems associated with central management such as co-located networks and configuration requirements, but it will allow the implementations of features such as power management, frequency hopping and roaming without compromising performance. The management function does not have to be stationary or pre-configured. It should be allowed to be dynamic to adapt to the needs of particular networks. This will ensure support for ad-hoc networks.

2) The standard should have the flexibility for expansion by including in the configuration fields the necessary information. Variable length periods for outbound, inbound and contention traffic would be defined here, as has been proposed in the R-TDMA protocol. Of course, the length of any of these fields could be set to zero in a particular implementation. This would allow a vendor to optimize their offering to satisfy a particular customer need.

3) The standard should require every node to understand and obey reservations such as those imposed by the R-TDMA protocol or the WHAT protocol. This will allow a network implementing a fully configured system supporting time-bounded and asynchronous services to coexist gracefully with a simple peer-to-peer network sharing a common geography. The reservation structure should be kept simple to allow for inexpensive state machine implementations.

4) Finally, a fully functional standard should be capable of supporting isochronous services in addition to asynchronous services. The central management agent would in this case perform all the scheduling functions necessary and would only use the non-contention periods for the time bounded traffic. For some PHY layers it may be necessary to limit the maximum time allotted to time bounded traffic to ensure enough bandwidth is available for asynchronous traffic (e.g. simple peer-to-peer collocated network is trying to share the media). This may not be necessary in a frequency hopping PHY where the two networks have to share the media only when their respective hops coincide.

CONCLUSION

Both centrally managed and peer-to-peer network topologies serve valuable purposes for certain applications. A robust and forward looking standard should allow for a range of wireless LAN implementations that can grow in complexity as application requirements change, while permitting inexpensive solutions to be commercially available. The proposals advanced to date have emphasized one type of network topology without adequately addressing the other. By combining the salient features of the two approaches, it is possible to devise a standard that meets the requirements of both simple, inexpensive designs and more complex, capable implementations alike. By creating a hierarchical feature set, simple networks can be made to interoperate with more complex networks without having to support all the services available to them. Such a standard is possible and bridges the current gap between proponents of both centrally managed and peer-to-peer approaches.