IEEE 802.11
Wireless Access Method and Physical Layer Specifications

Protocol Layering Alternatives for Practical Implementation

September 15, 1993

Michael Fischer
CHILD Systems, Inc.
PO Box 691086
San Antonio, TX 78269
voice (210) 614-4096
FAX (210) 614-8192

Abstract

The Reference Model currently defined for 802.11 addresses the architectural issues involved in supporting several, substantially different PHY types (DS, FH, IR) under a common MAC by subdividing PHY into a Medium Dependent Layer and a Medium Independent Layer, connected by a Convergence Layer. This approach to subdivision and convergence will require implementations of the resulting standard to be more complex than necessary. The added complexity results from the amount of communication across the MAC–PHY interface needed to permit MAC to support PHYs as diverse as the DS and FH PHYs proposed to date. This submission proposes changing the 802.11 Reference Model to make the "Medium Dependent Layer" (perhaps renamed as "Medium Adaptation Layer") a sub–layer within MAC, and to use the DTE/DCE interface at the top of the current PHY Convergence Layer as the MAC–PHY interface. This results in equivalent functional partitioning as with the current Reference Model, but requires substantially less ancillary information to be exchanged across the MAC–PHY interface in modular implementations of these two layers. The benefits of reducing this ancillary information exchange are the potential of reduced costs, reduced power consumption, improved modularity, and the facilitation of mixed–vendor station equipment.

1. Introduction

During the 802.11 efforts to date, the bulk of the work on the MAC layer has been concerned with the medium access and infrastructure aspects of MAC; while the bulk of the work on the PHY layer has been concerned with the wireless channel and the modulation techniques used to communicate over this channel. There has been relatively little detailed work concerning the manner in which a MAC implementation and a PHY implementation will interact with each other. This focus on the "top" of MAC and the "bottom" of PHY was appropriate for initial definitional work. However, continuance of this focus runs a major risk of creating a standard that (inadvertently) mandates implementations that are unnecessarily complex and inadequately modular. This risk originates with the subdivision of the PHY layer contained in the 802.11 Reference Model:

- MAC proposals, to the extent that they have addressed the MAC–PHY interface at all, have identified abstract services, without direct concern for what information is needed by the facility that provides the other side of this interface. In several cases, the necessary information is not information that a PHY implementation has directly available without violating conventional protocol layering or conveying additional (and otherwise unnecessary) information from MAC to PHY via a separate path from that used or normal information transfer.
- PHY proposals that discuss communication with the rest of the station, in addition to communication over the wireless link, generally stop at the "optionally exposed DTE/DCE interface" at the top of the Convergence Layer of the present Reference Model.

The result is that, as MAC definition progresses downward and PHY definition progresses upward, all the details that do not fall fully within MAC or fully within PHY accumulate in the "gray area" between MAC and PHY. Presently, this catch–all is the area labeled "Medium Independent Layer" which appears at the top of PHY in the present Reference Model diagram. As work proceeds toward resolving these accumulated issues, there is a choice to be made:

- Either continue to use the present Reference Model, with a "Medium Independent" subdivision of within the PHY layer. This provides excellent esthetics, encapsulating everything concerned with a particular wireless medium within PHY, but will require unnecessarily complex implementation at the MAC–PHY boundary if protocol layering is not to be violated.
- Or modify the Reference Model to place certain medium–specific functions in a subdivision within the MAC layer. This may be less esthetic, but is more practical because of the reduced complexity (hence reduced cost, power consumption, etc.) of the resulting network adapters.

The remainder of this document discusses the reasons for these differences in implementation complexity.

2. What is the "Medium Independent Layer?"

As is often the case in standardization, much of this problem concerns labeling. The confusion starts with the fact that the name "Medium Independent Layer" is only partially correct as a description of this layer —The layer is medium independent insofar as the DS, FH, or IR transceiver functions exist in the "Medium Dependent Layer;" however, the layer is not medium independent in its operation. In fact, each particular implementation of this "Medium Independent Layer" exists solely to provide the medium–specific functionality necessary to render the differences between the various 802.11 media transparent to MAC. In order to achieve the objective of a common MAC over several, diverse PHYs, as is specified in 802.11's PAR, these transparency functions are a significant undertaking, both to specify, to implement, and to validate.

1) It must be possible to implement of the "conventional" MAC functions, including access control, received frame validation, address filtering, interface with LLC, etc.; as well as the "wireless–specific" MAC functions including association, authentication, and channel sharing, implemented in a manner that is able to work with the widely–divergent characteristics of (at least) the Direct Sequence PHY and Frequency Hopping PHY, and (probably) the InfraRed PHY.

2) The differences in functional as well as operational characteristics of these PHY alternatives are considerably greater than between the various wired media alternatives supported by other 802 LANs. In particular, the concept of "medium" for the DS and FH PHYs includes not only the medium itself (frequency band) and the usage of that medium (modulation, spreading, radiated power level); but also various externally–visible characteristics of the usage of the medium (frequency hop sequencing and timing, fitting frames into remaining hop time, searching for access points, fragmenting and reassembling frames, assessing link quality, etc.). The "normalization" of these differences, which is the task of this "Medium Independent Layer," has not been of similar magnitude in any other 802 MAC–PHY interface.

   The physical layer definition process for wired LANs was generally done by selecting cables and signaling methods (or adopting methods already implemented) to provide the necessary capabilities under a pre–defined interface to the MAC. The approach of defining functionally–equivalent signaling schemes inherently yielded exposed MAC–PHY interfaces for wired LANs that were well–specified, both operationally and structurally. A significant benefit of these well–specified MAC–PHY interfaces was cost–effective support for a plurality of physical media in a manner that allowed interchangibility of PHY interfaces under a common MAC implementation.

3. An Alternative Layering Approach

It is certainly possible to define, and to implement, a "Medium Independent Layer" within PHY, however, it is undesirable to do so because, with the diverse characteristics of the current set of PHYs, this approach requires replication of functions and addition of information transfers simply to match the present Reference Model. By moving the "Medium Independent Layer" from the top of PHY to the bottom of MAC, these replicated functions and needless transfers can be avoided.

This change, which is illustrated in Figure 1, involves more re–labeling than functional change:

- The name "Medium Independent Layer" is shown as changing to "Physical Medium Adaptation (PMA) Layer." Since the relative position of this layer, between the service interface at the bottom of "MAC" and the DTE/DCE interface at the top of "Physical Medium Dependent," has not changed, the name could be left unchanged as well. However, the MAC is, by definition, medium independent, so the equivalent construct to "medium independent PHY" is "medium dependent MAC." The name change to "Physical Medium Adaptation" is proposed to avoid confusion with the pre–existing labeling, and to more accurately reflect this layer's functional role, which is adaptation between an abstract service interface and a concrete modem interface.

- The "Convergence Layer" is shown as being eliminated. Because the interface between the "Convergence Layer" and the "Medium Dependent Layer" was not an exposed interface, the sole purpose of making a distinction between these two layers is to permit an 802.11–compliant DTE/DCE interface to be added to a pre–existing wireless modem. With the "Medium Independent Layer" removed from PHY, the DTE/DCE interface becomes the MAC–PHY interface, which is inherently exposed. This appears to render the "Convergence Layer" unnecessary.
- The PHY is no longer subdivided, so the name "Medium Dependent Layer" can be dropped, and the physical layer can be referred to simply as "PHY."
- The MAC is now subdivided laterally, in addition to longitudinally for "MAC Management." Since the functions above the MAC service interface are unaffected by this change, there is no reason to re–label "MAC."

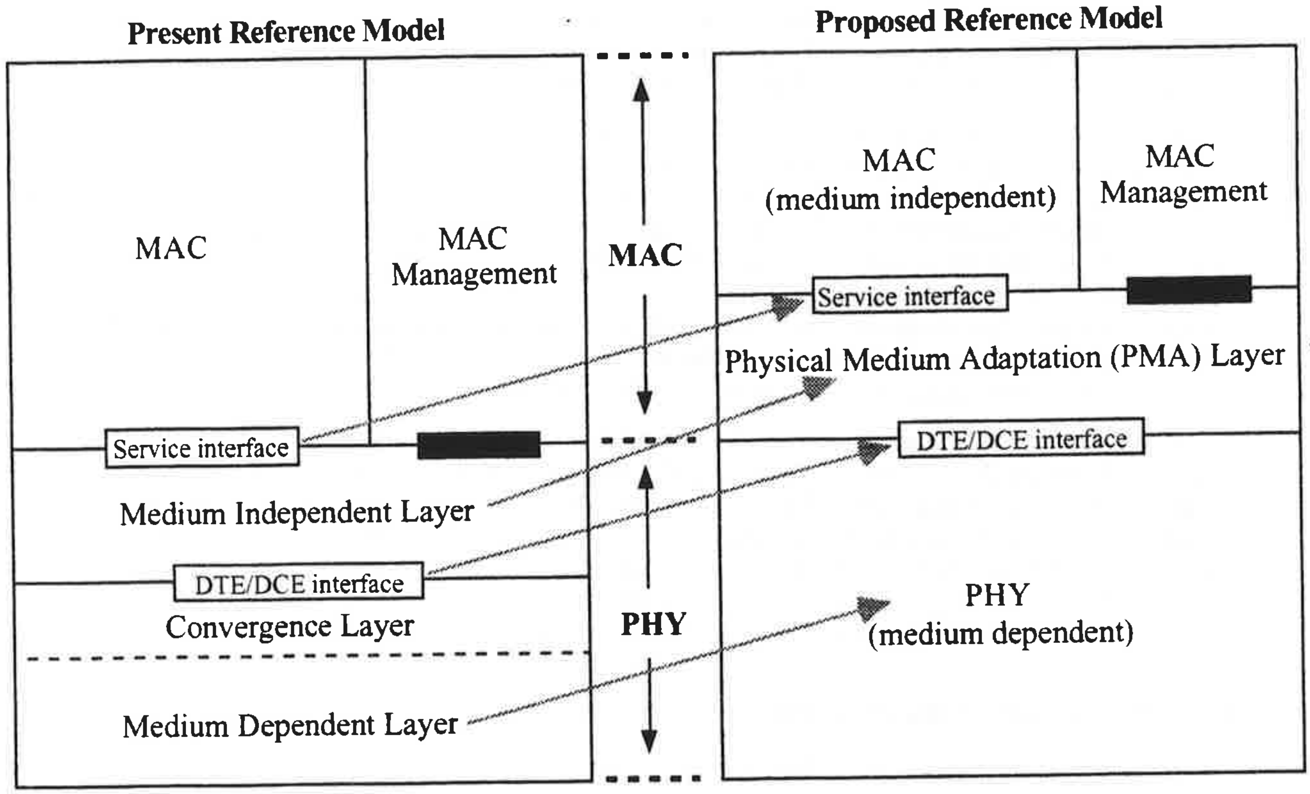

Figure 1. Reference Model Alternatives

4. Reasons for this Alternative Layering Approach

Wireless LANs introduce several situations that do not occur, or have not been of major concern, in wired 802 networks:

1) On wired LANs that support a plurality of media types, the differences between the various, media types relate primarily to cost, segment length, and data rate. Neither the communication control activities, the MAC access strategies, nor the error control strategies need to differ between the various media types. Accordingly, the physical layer functions for each of the media types can each be abstracted as functionally–identical, unintelligent interface elements that convey MSDUs to and from the particular physical wires (or fibers).

   On wireless LANs, the differences in among the various media types for characteristics such as range and data rate are negligible in comparison to the differences in communication control and MAC access requirements. Among the most significant of these, MAC–visible differences are:

   - A substantial portion of wireless stations are expected to need to operate on battery power. Power management involves control of power to the wireless transceiver on a time basis. Whether this control is performed in MAC or in PHY, it is necessary for this to be coordinated by MAC. In the absence of a Physical Medium Adaptation function within MAC, this requires a common, and synchronized, sense of time on both sides of the MAC–PHY boundary, as well as accurate timer facilities in both the MAC and the PHY implementations.
   - For the frequency hopping PHY it is necessary to control the timing and channel usage sequence of the hops. Whether this hop control is performed in MAC or in PHY, it is necessary for MAC to be aware of the hop timing, in order to avoid attempting to transmit a packet across a hop boundary as well as to maintain hop synchronization between the transmitting and receiving stations.
   - In the absence of a Physical Medium Adaptation function within MAC, the only way to provide the necessary information to MAC involves some sort of timing service across the MAC–PHY boundary. This increases the complexity of the PHY, to provide this interface, as well as of the MAC, to achieve sufficiently low–latency response to hop–timing indications from the PHY. The use of a common MAC in this manner requires this complexity, and the added overhead of supporting this real–time information exchange is incurred even when using non–hopping PHYs.

     Accordingly, the objective of a single, practical MAC is seriously jeopardized by the absence of PHY–specific adaptation facilities as a subdivision of MAC.

2) The headers of PHY frames on wired LANs are preambles containing solely the signal patterns needed for bit and byte synchronization, frame delimitation, and (for some networks) carrier detection. The remainder of the wired LAN PHY frame is the MSDU, so a very limited set of MAC–PHY interface services, on the general level of "send bit" and "receive bit" are sufficient.

   The headers of PHY frames on wireless LANs include all of the features of wired LAN preambles, plus generally requiring additional signal pattern content and ancillary information for functions such as antenna diversity processing, despreading, unique delimiters to reliably distinguish between end–of–frame and loss–of–signal, identification of payload bit rate in variable bit–rate systems, etc. Inter–station communication control also requires conveyance of PHY–related information, such as hop synchronization and logical network identification, between stations. This information can be conveyed in a PHY header or in the MAC header:

   - If the information is placed in the PHY header, significant redundancy in required, because this information is not part of the FCS–protected payload. This redundancy consumes transfer bandwidth, which is the scarcest resource on wireless networks. The processing of this redundantly–encoded information adds complexity to the PHY implementation.

- If the information is placed in the MAC header, and the present Reference Model is maintained, the relevant header information must extracted by MAC and returned to the "Physical Medium Independent" portion of PHY for interpretation after the MSDU has already been transferred from PHY to MAC and the FCS validated. It is violation of protocol layering for PHY to directly interpret the contents of an MSDU, and is also dangerous, because the contents of an MSDU are not known to be valid until the FCS has been checked. Accordingly, processing this PHY–specific information below the MAC–PHY boundary requires replication of processing capability on both sides of this boundary, but also requires a separate, concurrent transfer path for extracted header information from MAC back to PHY. [NOTE: A mechanism (kludge?) of this type is used in 10broad36, which was a *post–facto* attempt to utilize the 802.3 MAC layer over preexisting broadband wired media. There is no apparent reason to mandate this type of complexity for 802.11 when there is no preexisting signaling standard nor installed base.]

3) On wired LANs, it is axiomatic that signals received by the PHY from the physical medium pertain to conveyance of an MSDU of relevance to the associated MAC — signals from other networks are on other cables! The only plausible interpretations of incoming signals from wired LAN media are as MSDUs or as noise. Accordingly, all data bits extractable from incoming signals by a PHY layer of a wired LAN are passed to the MAC for possible interpretation.

   On wireless LANs, the physical medium does not provide this inherent isolation of physical networks. Signals arriving received by a wireless PHY from the physical medium may convey

   - MSDUs from this network,
   - MSDUs from another network that is using the same PHY,
   - interference from another network that is using a different (802.11–compliant) PHY,
   - interference from other users of the same frequency band, or
   - noise.

   Therefore, the PHY layer of a wireless LAN may need to filter certain non–noise receptions, and may need to pass information to MAC pertaining to receptions that can be demodulated by the PHY, but which are not MSDUs of this network. The forms and discrimination of these signal types is not identical or all PHYs, but some of the distinctions, such as logical network identification, may not be reliably accomplished without examining the contents of the MSDU. This is another situation where the function is straightforward to implement in a Physical Medium Adaptation Layer of MAC, but would require information from a received MSDU to be passed back from MAC to "Physical Medium Independent" PHY if there is no such facility within MAC.

5. Benefits of this Alternative Layering Approach

By adopting the alternative Reference Model discussed herein and illustrated in Figure 1, the following benefits are obtained:

1) Information pertaining to the particular physical link in use can be exchanged between stations in MSDUs, and interpreted in a PHY–specific manner by the Physical Medium Adaptation layer within MAC. This permits exchange of reliable (FCS protected) PHY–related information without creation of an extra path to return this information from MAC to PHY for processing after frame validation within MAC.

2) All of the decisions pertaining to PHY interface control, including power management, hop timing (both for hop synchronization and for coordination of MAC frame transfer with hop time availability), channel selection, and scanning for access points can be accomplished within MAC. This permits simpler, lower cost implementations because MAC–aware, real–time functionality does not have to replicated on both sides of the MAC–PHY interface.

2) The complexity of the MAC–PHY interface is reduced to that of a simple, "unintelligent" DTE/DCE type interface. This both functional and electrical uniformity at this interface, thereby permitting replacement or substitution of PHY modules.

6. Issues Addressed

1.5 Should the protocol model generated during the July 1992 meeting be adopted by 802.11?

NO. There should be a medium–dependent subdivision added to MAC (suggested name is "Physical Medium Adaptation Layer"), and the subdivision on PHY should be eliminated. This is discussed in further detail in this document.

9.3 Must the same MAC work in a minimum system and a maximum system?

YES.

12.1 What is the MAC–PHY interface?

The MAC–PHY interface is generally a DTE/DCE interface of the type discussed previously for the "optionally exposed interface" between the "Convergence Layer" and the "Medium Independent Layer" within PHY. The functions at this interface include

- serial transmit and receive bit streams, using clocks provided by PHY;
- direct control and status signals for transmitter and receiver enable, clear to send, carrier sense, and receive data available;
- serial command transfers from MAC to PHY for functions such as setting power levels, setting transmit bit rate, setting receive thresholds and acceptable quality levels, defining LBT deferrals, selecting frequencies, and requesting a status transfer; and
- serial status transfers from PHY to MAC, pursuant to request, for information such as received signal quality, received bit rate, and PHY–specific parameter values.

12.3 What is the intelligence level of the MAC–PHY interface?

The MAC–PHY interface is an "unintelligent" interface, permitting a single, "intelligent" MAC, with a replaceable, PHY adaptation function, to directly attach, both logically and electrically, to a plurality of different PHY types.

12.4 Is the layer that provides PHY independence the same as the MAC–PHY interface?

NO. This independence is implemented in a subdivision of MAC dedicated to PHY adaptation. This permits processing of the relevant information from received frames, after MSDU acceptance and validation by the receiving MAC, without an extra mechanism to return this data to PHY for processing.

12.8 Does a PHY independence layer need to specified in the MAC?

YES.

12.9 Should data and control information be passed simultaneously across the MAC–PHY logical interface?

YES.

18.2 Will the standard support one MAC driving multiple PHYs of different rates?

YES. The use of a PHY adaptation layer at the bottom of MAC allows such multiple–PHY support, provided that the necessary parameters regarding the PHY capabilities can be requested by this PHY adaptation layer via the MAC–PHY interface.

24.3 How will multiple PHY support for the MAC be specified?

The PHY adaptation layer at the bottom of MAC processes PHY–specific information, and inserts and extracts such information to/from MAC headers being exchanged over the wireless media.

24.6 Does the PHY layer provide the PHY type to the MAC layer?

YES.

24.7 Will the MAC standard specify the support of multiple PHYs?

YES.

24.8 What functions are required in the Medium Independent PHY layer?

NONE. This layer is unnecessary. All medium–independent functions are able to be performed in MAC. The "adaptation" or "convergence" function is needed at the bottom of MAC, not the top of PHY. The appropriate MAC–PHY interface is to the "medium dependent" portion of PHY.

24.9 Given a Frequency Hopping PHY, which protocol entity is responsible for the real time aspect of the PHY?

The PHY adaptation layer within MAC controls the timing of the hopping and the channel sequence. The (medium dependent portion of) PHY controls the tuning changes necessary to execute the hops commanded from the PHY adaptation layer.