IEEE 802.11
Wireless Access Method and Physical Specification

DFWMAC

Distributed Foundation Wireless Medium Access Control

Wim Diepstraten
NCR WCND-Utrecht
NCR/AT&T Network Product Group
Nieuwegein The Netherlands
Tel: (31)-3402-76482
Fax: (31)-3402-39125
Email: Wim.Diepstraten@utrecht.ncr.com

Greg Ennis
Symbol Technologies
1101 S. Winchester Blvd., Suite B-110
San Jose, CA USA 95128
Tel: (408) 249-9890
Fax: (408) 446-4630
Email: gennis@aol.com

Phil Belanger
Xircom
2105 Landings Drive
Mountain View, CA USA 94043
Tel: (415) 691-2500
Fax: (415) 691-1064
Email: pbelange@xircom.com

Abstract

This document presents a proposal to IEEE 802.11 for a wireless medium access control protocol. Most of the basic concepts within this document have been presented in several prior submissions to IEEE, known by their acronyms **WMAC** (NCR and Symbol) and **WHAT** (Xircom). The protocol incorporates a simple, distributed coordination function based upon CSMA with collision avoidance, together with an optional point coordination function which provides for contention-free transmissions supporting time-bounded services.

Table of Contents

1. Introduction

1.1 Purpose and Scope of this Document

This proposed DFWMAC protocol is part of a larger protocol architecture, which when implemented within actual equipment will provide a complete wireless LAN system. Two types of wireless LANs are intended to be supported by the equipment which implements these protocols:

- Ad Hoc Wireless LAN, which is a (typically small) network of stations, normally all within range of each other station's wireless transceivers.

- Infrastructure Wireless LAN, which includes special nodes called *access points* which are interconnected via a (typically wired) *distribution system* to provide a network for wireless stations encompassing a greater area than would be possible in the ad hoc case, and which provides access to existing wired networks.

The stations within a wireless LAN may be mobile. In the case of an infrastructure LAN with multiple access points, the protocol provides for the movement of stations from the area supported by one access point to that supported by another.

Another provision of the DFWMAC protocol described here is support for power conservation within the stations. Since mobile stations are likely operating on lightweight batteries, power conservation is a critical feature for these stations.

As is indicated by the protocol's name, the basic access method is distributed, based upon CSMA with collision avoidance. Built on top of this foundation is a mode of operation allowing for ***contention-free*** service. In this optional mode of operation, the time domain is formatted into periodic ***superframes*** consisting of alternating contention-free and contention periods. The superframe mechanism and the resulting contention-free period are supported optionally only in certain infrastructure configurations.

Both ***asynchronous data service*** and ***time bounded services*** are supported. Asynchronous services can be supported using the basic contention access method, which is supported by all stations and access points. In configurations implementing the optional contention-free mode, both time bounded service and contention-free asynchronous service can be supported.

Section 2 of this document provides a general overview of DFWMAC. This is followed in Section 3 with a detailed description of both the distributed and point-coordinated access methods. Section 4 covers DFWMAC's synchronization mechanisms, which are used to support power management as described in Section 5, and which also support superframe timing and PHY synchronization. Section 6 describes the manner in which both time-bounded services and asynchronous contention-free services can be supported within the superframe structure. Section 7 presents the DFWMAC frame formats.

1.2 Background

This document presents a proposal to IEEE 802.11 for a wireless medium access control protocol. Most of the basic concepts within this document have been presented in several prior submissions to IEEE, known by their acronyms **WMAC** (NCR and Symbol) and **WHAT** (Xircom). Although there were some differences between those two proposals, they bore a strong "family resemblance", including a distributed CSMA/CA foundation into which contention-free and time-bounded services can be incorporated. The present proposal is a merger of the two, selecting the best ideas from each and developing an integrated approach. During the development of this combined protocol approach it became apparent that in certain areas it was possible to incorporate refinements and improved techniques not present in either of the source proposals, using ideas which had been previously discussed not only among the authors but within the committee as a whole.

For those familiar with the two earlier proposals, the following summarizes the relationship between the present proposal and the earlier ones:

- The basic distributed access method is **CSMA with collision avoidance**, as was true in both of the earlier proposals. The overall system performance has been improved via a parameterized, per-frame approach to the use of RTS/CTS (as in the WHAT proposal), with the goal of minimizing both collisions and overhead. Frames are transmitted either with or without RTS/CTS, depending upon the parameterization. This approach allows a system to be parameterized so as to use RTS/CTS when it is most beneficial while eliminating overhead for other transmissions. The resulting system performance is thereby enhanced by using the optimal transmission approach for each frame.

- The actual collision avoidance algorithm proposed here is a refinement of that given in the two earlier proposals.

- Synchronization is achieved in the infrastructure case using the WMAC mechanism. A new distributed approach to synchronization in the ad hoc case has been added, including provisions for hopper synchronization.

- The power management provisions of WMAC have been maintained, both for the infrastructure and the ad hoc cases.

- The WHAT **MPDU ID** concept has been incorporated into all multiframe exchanges so as to minimize overhead.

- The support of time-bounded service involves a **point coordination function** together with a **superframe** concept in which contention-free access to the medium can be guaranteed. This approach is taken essentially from the WMAC proposal, improved through the use of the WHAT "net allocation vector" concept.

- The contention-free period within the point-coordination superframe is now able to support asynchronous as well as time-bounded services.

- Updated frame formats are defined, providing common labels for functions which were present in both of the previous proposals (for example, the functions of the WMAC **TIM** frames and the WHAT **Announce** frames are subsumed under the **Beacon** frames within the present protocol).

- Material has been included regarding association and re-association, the basic concepts of which were implicit in the WMAC and WHAT approaches though not described in the earlier proposal documents.

The remainder of this document will make no further references to specific WMAC or WHAT mechanisms, focusing instead on a detailed description of the integrated protocol. In the interest of providing a complete document, certain sections contain material replicated from the earlier papers.

1.3 Acknowledgments

In addition to the authors of this paper, the following individuals have also contributed to the development of the concepts in this protocol: Ken Biba and Bill Baugh (Xircom), LaRoy Tymes, Sarosh Vesuna and Chris Zegelin (Symbol), and Jan Kruys, Henri Moelard and Henk van Bokhorst (NCR).

The authors also benefited from numerous discussions with Jim Schuessler (National Semiconductor) regarding time-bounded and contention-free services.

1.4 References

[1] IEEE P802.11-93/70 A Distributed Access Protocol Supporting Time-Bounded Services, Wim Diepstraten, NCR

[2] IEEE P802.11-93/40 The Wireless Hybrid Asynchronous Time-Bounded MAC Protocol, Phil Belanger and Ken Biba, Xircom

[3] IEEE P802.11-93/95 Distributed Access WMAC, Synchronization and Power Management Mechanisms, Greg Ennis, Symbol Technologies and Wim Diepstraten, NCR

[4] IEEE P802.11-93/159, Frame Priority in CSMA/CA MAC Protocol, Rick White, Motorola

2. General Overview

2.1 Basic and Extended Service Area

The basic building block of a wireless network is a cell called a BSA (Basic Service Area). A group of stations within a BSA that are associated with each other to communicate, forms a BSS or Basic Service Set. The size of a cell is dependent on the environment and the characteristics of the wireless transceivers. Larger areas can be covered by using multiple BSAs, which can be interconnected using Access Points (APs) and a Distribution System. This is called an ESA or Extended Service Area. The set of stations within multiple BSSs which are connected via a Distribution System is an Extended Service Set (ESS). Each ESS is distinguished by an ESS-ID, while within one ESS, each BSS is distinguished by a BSS-ID. These IDs together form a NetworkID.

2.2 Ad-Hoc Network

A special case is a single ESA that does not have an Infrastructure. This is called an Ad-Hoc Network, which can be formed by a number of stations (such as in a meeting room, for instance). This network should be able to be set up fast, and would perhaps only last for the duration of the meeting, without the need for any infrastructure provisions.

It would further allow communication between stations without the need to be registered to an available infrastructure, and is independent of any installed infrastructure.

The following figure depicts an ad hoc network.

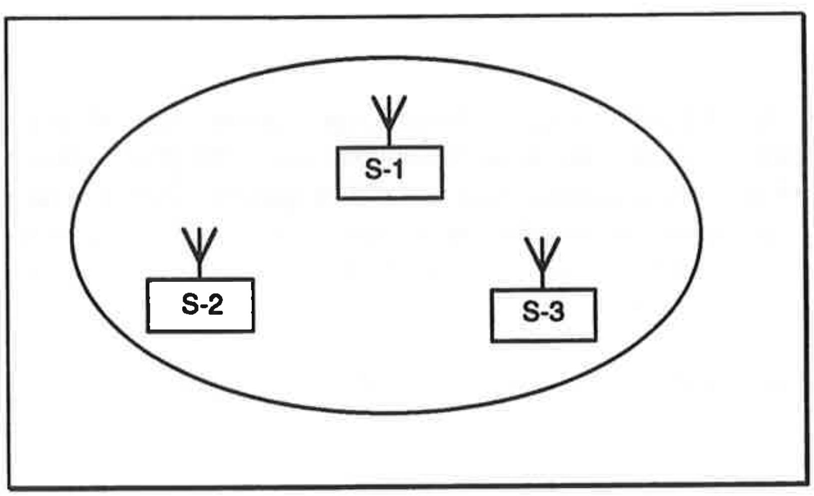

Figure 2-1: Ad Hoc Network

2.3 Infrastructure Network

To interconnect different BSS's to form a larger coverage ESS, an infrastructure is needed, which consists of specialized stations called Access Points (AP) and a Distribution System to interconnect the Access Points.

2.3.1 Access Points

An Access Point is part of the infrastructure and provides the following basic functionality:

- **Authentication, Association and Reassociation** services that allow stations to remain connected to the infrastructure while moving from one BSA into another.

- **Power Management** functions that allow stations to operate in a Power Save mode. This includes a temporary frame buffering function, and associated traffic announcement functionality to allow station transceivers to be powered down most of the time.

- **Synchronization** functions which ensure that all stations currently associated with the AP are synchronized on a common clock. This synchronization is used, for example, in supporting time bounded services, hopper management, and power management.

2.3.2. Distribution System

To interconnect different BSS's via AP's, a Distribution System is needed. It also provides internetworking facilities to interface to an existing network. The Distribution System can be built from wired network components, or it could be (partly) wireless. The minimum Infrastructure configuration consists of a single Access Point that is not connected to other AP's or to a standard wired LAN by a distribution system. Stations within one BSS are associated with a specific Access Point.

Figure 2-2 below depicts the various components of a typical infrastructure network.

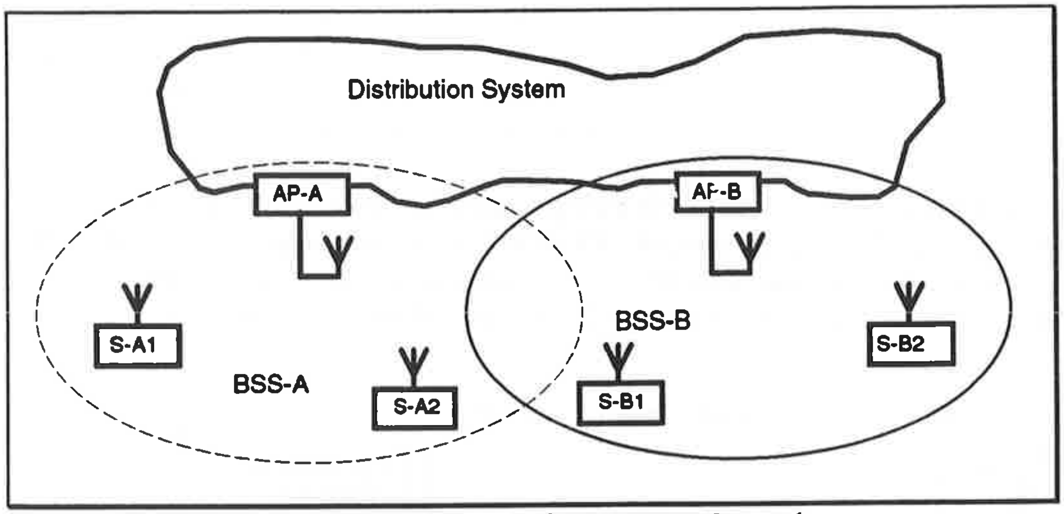

Figure 2-2: Infrastructure Network

2.4 Basic DFWMAC Protocol Approach

DFWMAC incorporates a basic distributed access method and an optional centralized access method. These two access methods can coexist, and in fact the distributed access method forms the foundation for the centralized method. These two methods (or "coordination functions") are then used to support two types of traffic: *asynchronous* and *time bounded.*

2.4.1 Distributed and Point Coordination Functions

A **coordination function** in a wireless LAN is a mechanism which determines when a particular station is allowed to transmit. With a **distributed coordination function**, this determination is the responsibility of the individual nodes and may result in multiple simultaneous transmissions, while a **point coordination function** centralizes this decision at any given moment and consequently ensures that only a single node may transmit.

The fundamental access method within DFWMAC is a distributed coordination function known as CSMA with collision avoidance, or **CSMA/CA**. The use of this access method is mandatory for all stations and access points, and is used both within ad hoc and infrastructure configurations.

A station wishing to transmit under the CSMA/CA rules operates essentially as follows: the medium is sensed to determine if another station is transmitting. If not, the transmission may proceed, while if the medium is sensed busy the station must defer until the end of the current transmission. After deferral, the station will select a random interval and will check that the medium remains idle for that interval - this is the collision avoidance aspect of the method. A refinement of the method can be used under various circumstances to further minimize collisions - here the transmitting and receiving station exchange short control frames (**RTS** and **CTS** frames) prior to the data transmission. The details of CSMA/CA and RTS/CTS exchanges are presented in Section 3.

Since multiple frames are often involved in a single exchange, the concept of a MAC Protocol Data Unit (MPDU) is introduced. An MDPU may, for example, consist of RTS, CTS, Data and Ack frames. The basic flow of a non-broadcast (i.e. a "directed") CSMA/CA data transfer (with or without RTS/CTS) is depicted in the following figure.

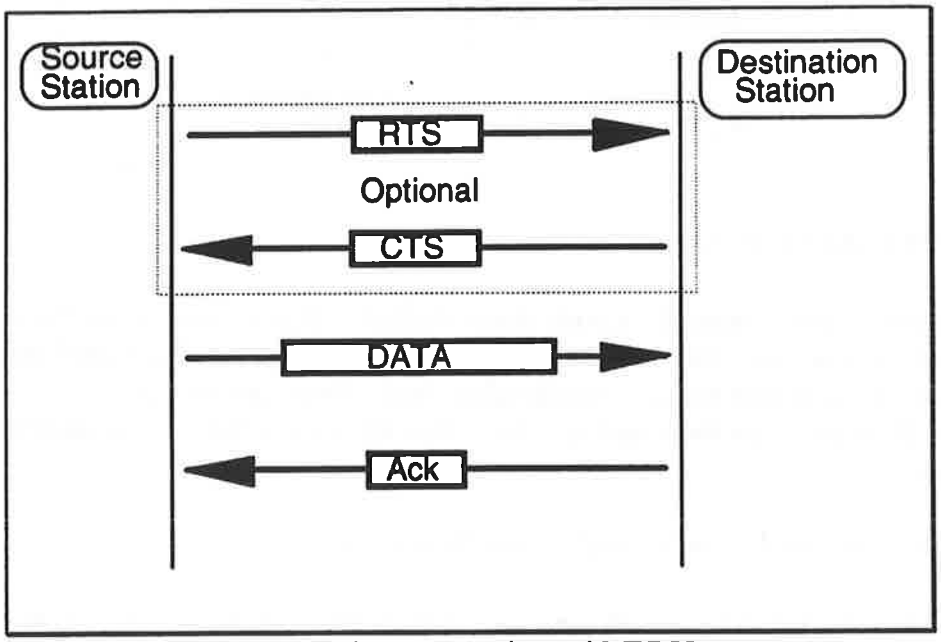

Figure 2-3: Directed MPDU

DFWMAC also incorporates a point coordination function, implemented on top of the basic CSMA/CA protocol. This access method uses a point coordinator (typically the access point) to determine which station currently has the right to transmit. The operation is essentially that of polling with the point coordinator playing the role of the polling master.

NOTE: the point coordination function of DFWMAC could also be viewed as a specialized token passing system, in which the token is forced to circulate through a special node (the point coordinator) which is guaranteed to be within range of all other nodes. A standard token passing approach in which the token is passed directly between stations is impractical in a wireless system due to the fact that not every station is necessarily able to hear every other station. Use of the access point as the point coordinator is effective since it is likely that a high proportion of traffic will pass through it.

The support of the DFWMAC point coordination function is optional, and in fact requires that the network configuration involves no overlapping point-coordinated BSS's on the same channel (a BSS using the distributed coordination function can, however, overlap with a similar BSS or even with a single point-coordinated BSS).

The point-coordination function is built up from the basic CSMA/CA function through use of an access priority mechanism. The CSMA/CA distributed algorithm mandates that a gap of a specified duration exist between subsequent frames (an **interframe space**, or **IFS**). A transmitting station must ensure that the medium is idle for the appropriate IFS duration before attempting to transmit. Different classes of traffic can be defined through the use of different values for IFS, thereby creating prioritized access to the medium for those classes with a shorter IFS. DFWMAC defines several such IFS's, including a Point IFS (PIFS) and a Distributed IFS (DIFS) for data frames transmitted via the point and distributed coordination functions. Since PIFS is shorter than DIFS, point-coordinated traffic essentially has priority access to the medium.

The access priority enjoyed by point-coordinated traffic is utilized to create a **contention-free** access method. The priority access of the PIFS allows the point coordinator to "seize control" of the medium away from the other stations. The point coordinator can then control the frame transmissions of the stations so as to eliminate contention. Also used in this regard is a mechanism in which the point coordinator can signal to the other stations that such a burst is occurring, causing them to be silent during that period even if they don't directly hear the traffic (this is essentially the **net allocation vector**, or **NAV**, concept discussed in Section 3).

The two access methods are integrated in a **superframe** concept in which a contention-free burst occurs at the beginning of the superframe, followed by a contention period. This is also described in Section 3. The following figure depicts this mode of operation.

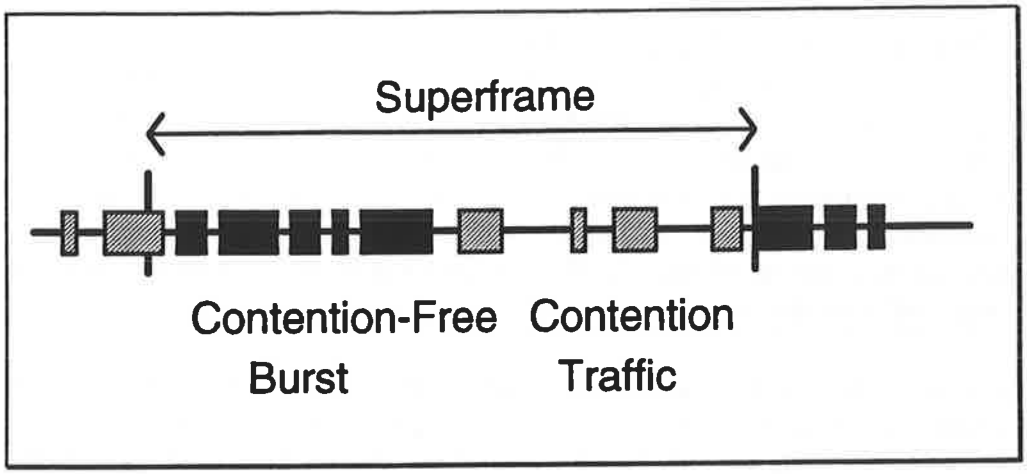

Figure 2-4: Superframe Structure

2.4.2 Asynchronous Service

Asynchronous service is provided within both ad hoc and infrastructure configurations, using both the CSMA/CA and point-coordination functions.

DFWMAC enhances the robustness of CSMA/CA for asynchronous traffic through the use of a *MAC level Acknowledgment*. The CSMA/CA part is designed to share the medium in a spectrally efficient way. The MAC-level Acknowledgment allows the detection of a lost frame and its recovery by retransmission at the MAC level. This is the way that asynchronous traffic is managed within a pure DCF configuration (e.g. an ad hoc network or an infrastructure network supporting only the DCF). Similarly, CSMA/CA + Ack is used within the contention period of a superframe when a point coordination function is present. All Asynchronous frame transmissions use the CSMA/CA+Ack approach, except for Broadcast/Multicast frames directed to stations, which are not acknowledged. Broadcast/multicast frames which are generated by stations and are to be forwarded by an access point will be acknowledged by the AP.

To enhance the robustness of the basic CSMA mechanism in "hidden station" situations, a particular frame may be transmitted using an RTS/CTS mechanism in which the sender and receiver declare their intent to transmit for a specific duration to other stations. Since this mechanism is particularly appropriate for protecting long frame transmissions and transmissions from certain nodes, its use is governed by a per-station parameter which determines which frames (by size) are to be transmitted with RTS/CTS and which are to be transmitted without RTS/CTS.

The details of the proposed CSMA/CA mechanism can be found in Section 3 of this document.

Contention-free asynchronous transmissions can be handled by the point coordination function. Here the point coordinator may allow station to participate in the contention-free burst, during which it transmits its asynchronous frames. A station participating in such a service can also use the contention period for its asynchronous frames, and the contention-free transmission opportunities can be viewed as a performance booster for the given station. The use of the contention-free period for asynchronous transmissions is possible within a superframe only if the time bounded stations have already had their requirements met.

2.4.3 Time Bounded Service

Stations requiring a bounded delay variance in their transmissions can use the DFWMAC time bounded service. This is implemented within the contention-free (point-coordinated) portion of a superframe. Each such station will establish a time-bounded connection, with a parameter defining the frequency with which it requires a transmission opportunity. The point coordinator (typically the access point) will ensure that the station receives such opportunities at regular intervals (e.g. once within each superframe period).

Since the time bounded service is based upon the point coordination function, it is an optional service and is possible only in configurations in which the point coordinated BSS overlaps no other point coordinated BSS.

The Acknowledgment mechanism implementation is somewhat different for the Time Bounded Service.

The details of Time Bounded Service are discussed in Section 6.

2.5 Time Coordination

A critical function within DFWMAC is the coordination of time among all stations within a BSS. This is used for the superframe concept as discussed above, but is also used for other important functions such as power management (discussed below) and PHY synchronization (particularly for FHSS systems).

In the infrastructure case, time synchronization is achieved through the use of the access point as the timing master. The access point periodically generates **beacons** which contain the current clock value, which is then copied by the stations. A station need not hear all beacons reliably, since they are used only to calibrate the station's own timing clock.

In the ad hoc case, the stations share the beacon generation responsibility among themselves in a distributed fashion. An algorithm is defined by means of which stations which were previously not synchronized may converge to a common clock value. Once synchronized, the common time reference can be used for power management and hop timing.

Synchronization is discussed in detail in Section 4.

2.6 Power Management

An essential part of the DFWMAC is the ability to conserve mobile station power consumption as much as possible. In Infrastructure mode this is achieved by the following principle steps:

- Synchronize all stations operating in Power Save Mode within one BSS.
- Allow stations to power down their transceivers for most of the time.
- Temporarily buffer incoming frames for sleeping stations in an Access Point.
- Announce the buffered frames to the stations at regular predetermined intervals, so that a station can sleep until it wishes to hear such an announcment, at which point it can take steps to ensure that any buffered frames are properly received.

Different Power Management modes are provided to trade-off station power consumption and performance.

In Ad-Hoc Networks a similar approach is followed that makes use of the fact that all stations are synchronized with each other. No central buffering function is available in this configuration, so a frame may be buffered within a station until all transceivers are awake.

2.7 Scanning and Station Mobility

Scanning is a process in which a station examines its immediate environment to determine a BSS with which to associate. This is typically performed both at initialization and periodically during normal operation. In a multi-channel system this involves scanning the various channels.

In the infrastructure case, stations will normally "Associate" with an Access Point as part of their initialization process. By scanning they will determine the "best quality" link by listening to frames which are transmitted by each Access Point. The signal level and quality of the frame reception from each of the AP's that are received during a certain scan time interval can be evaluated by the stations. The stations may elicit the generation of such a frame through the transmission of a **probe**. Stations can then Associate with the new Access Point. The Access Point will in turn follow a hand-off protocol across the distribution system, to inform the other Access Points regarding the new association.

In the ad hoc case, scanning is required to identify the PHY characteristics of the desired BSS (e.g. frequency channel and/or hopping sequence).

2.8 PHY Independence

DFWMAC has been designed specifically to accommodate a wide variety of PHYs.

2.8.1 FHSS, DSSS, and Infrared

The three primary physical layers identified for IEEE 802.11 include Frequency Hopped Spread Spectrum, Direct Sequence Spread Spectrum, and Infrared. DFWMAC can be implemented using any of these PHYs. Specific support has been included to consider the different scanning requirements of the different PHYs. The synchronization provisions of DFWMAC also are particularly useful in the case of FHSS.

Certain operations of DFWMAC will differ depending upon the PHY. This can be implemented as a "PHY-dependent submachine" within the general DFWMAC state machine.

2.8.2 Single Channel and Multi-channel PHYs

DFWMAC accommodates both single channel and multiple channel PHYs. Single channel PHYs can support multiple DFWMAC BSS's using the basic CSMA/CA distributed coordination function. The use of a single point-coordinated BSS is possible in a single channel PHY (with, perhaps, multiple DCF BSS's), with multiple point-coordinated BSS's possible only if spatial isolation is guaranteed.

Multi-channel PHYs are supported through the explicit provision of multi-channel scanning algorithms, both in the infrastructure and ad hoc environments.

2.8.3 Multiple Bitrates

DFWMAC has been designed with the intent of allowing the use of multiple bitrates simultaneous within a single BSS. In general, control information which needs to be received by multiple stations is transmitted in separate control frames, allowing for transmission of this information at a common bitrate implemented within all stations. Also, duration (rather than length) has been used within the frame formats so as to provided a bitrate-independent representation of frame sizes whenever this was deemed necessary. This document does not cover all aspects of this issue, which will be addressed in further contributions.

2.8.4 Other Bands

Although the initial focus within IEEE 802.11 for radio has been the 2.4 GHz ISM band, it is a goal of DFWMAC that it be adaptable in other radio PHYs in other bands, such as the 1.9 GHz PCS and ETSI STC/RES10 HIPERLAN 5.2 GHz bands. Support of PHYs specific to various international environments is also a goal.

2.9 Encryption and Compression

The present document does not specify actual encryption or compression mechanisms. However, the frame format includes control fields which identify whether or not a particular frame is encrypted or compressed. These fields would allow the MAC entity to take proper action on a per-frame basis. It may be that multi-bit control fields are required rather than the single bits in the current proposal.

The exact manner in which encryption (based upon 802.10) or compression should be incorporated into the MAC is an ongoing topic of research within the 802.11 committee.

3. Access Method

DFWMAC supports two access methods: a mandatory distributed coordination function method which is available in both ad hoc and infrastructure configurations, and an optional point-coordinated coordination function which is available within certain infrastructure environments and which can provide time-bounded services.

3.1 Basic Distributed Coordination Function (CSMA/CA)

The basic medium access protocol is a Distributed Coordination Function (DCF) that allows for automatic medium sharing between similar and dissimilar systems through the use of CSMA/CA (Carrier Sense Multiple Access with Collision Avoidance).

The access method uses a carrier sense mechanism in which the PHY detects whether signal energy in the occupied band is above a threshold, so as to determine whether the wireless medium is available for that station. To reduce the probability for access collisions, and to provide fair access to the medium by all stations, a random backoff mechanism is used to randomize the gap between stations accessing the medium.

This scheme is further enhanced by a MAC level provision that maintains a "Virtual CS" indication called the **Net Allocation Vector** or **NAV**, based on duration information that is being announced in special RTS/CTS frames prior to the actual exchange of the data.

The CSMA/CA protocol is designed to reduce the collision probability between multiple stations accessing a medium, at the point where they would most likely occur. This is just after the medium becomes free, following a busy medium as indicated by a CS function or the NAV, because multiple stations could have been waiting for the medium to become available again. The probability that multiple stations using CSMA would access the medium at exactly the same time is high immediately following a medium busy situation caused by a transmitted frame. This is the situation where a random backoff arrangement is needed to resolve medium contention conflicts.

3.1.1 MAC-Level Acknowledgments

MAC level Acknowledgment is an inherent part of the access protocol. This is achieved by the CSMA/CA + Ack access scheme. To allow detection of a lost frame (due to interference or collisions), an Ack is returned by the destination station immediately following a successfully received frame. The gap between the received frame and the Ack frame is such that it has priority over access of the medium by all other stations which are waiting for the medium to become available.

The Ack is transmitted by the (addressed) receiving station only when the CRC of the received frame is found correct, so it is a positive Acknowledge mechanism. The lack of a returned Ack can be used by a MAC transmitter to recover from this error by a retransmission of the frame after a random "Retransmission-Backoff".

3.1.2 CSMA/CA Access Procedure

The basic access mechanism is shown in the following timing diagram.

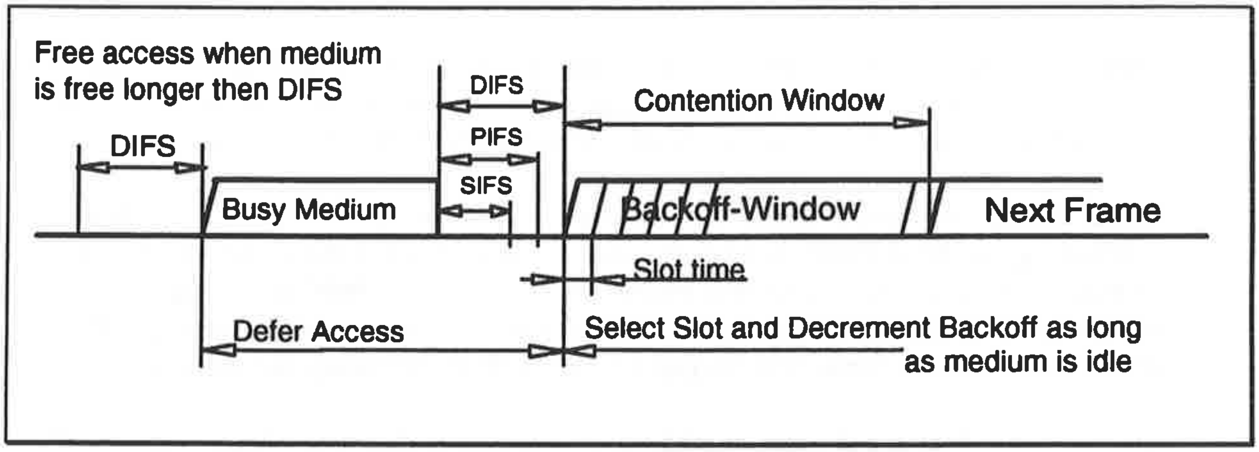

Figure 3-1: CSMA/CA Access Method with PCF Priority

The key procedure is that a station that wants to access the medium needs to sense the medium first to ensure that a particular minimum silence period (IFS or Inter Frame Space) has elapsed, before the medium is accessed. Three different type of access priorities are distinguished, using different values for the minimum silence period.

Short Priority:

This priority level is used for all immediate response actions. This priority levels used for an Ack frame, immediately following a received frame, a CTS frame immediately following a received RTS frame, and by a station responding to any polling as is used by the Point Coordination Function (PCF). The corresponding interframe space is called the Short IFS, or SIFS.

PCF Priority:

This priority level is used by the PCF in the AP to send any of the Contention Free Period (CFP) frames. The AP will send the next-in-line queued CFP frame after it finds the medium free for the period PIFS (PCF Interframe Space), during a CFP-Burst .

DCF Priority:

The DCF priority is used by the Distributed Coordination Function to transmit Asynchronous frames in the Contention Period. Asynchronous stations that want to start transmission of an RTS frame, or a DATA frame (without the RTS/CTS option), will monitor the medium for at least a time DIFS (DCF Inter Frame Space) of silence, after the NAV and the CS function indicates a free medium. When the medium was found busy, then the DFWMAC will defer until an DIFS is detected, and then a random access backoff procedure is started.

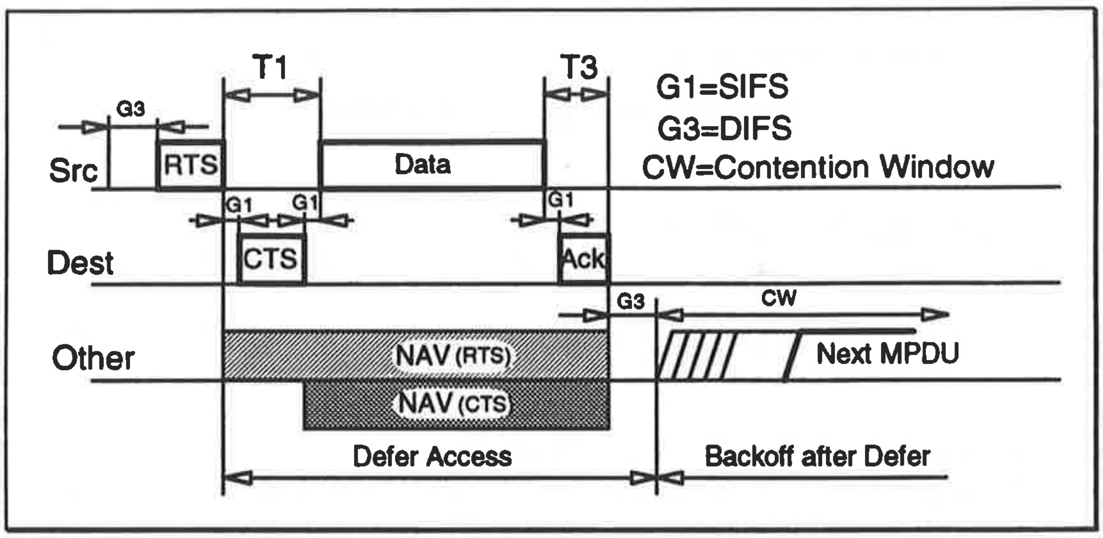

Figure 3-2: RTS/CTS/DATA/Ack MPDU

The above figure shows the Directed MPDU format that is used in the Contention Period, with the use of the RTS/CTS option. The RTS and CTS frames contain a "Duration" field that indicates the medium occupancy time of the MPDU from the end of the RTS or CTS frame until the end of the Ack frame. Non-addressed stations need to interpret these frames, and maintain the Net Allocation Vector (NAV). The figure indicates the NAV for stations that can hear the RTS frame, while other stations may only receive the CTS frame, resulting in a NAV as shown.

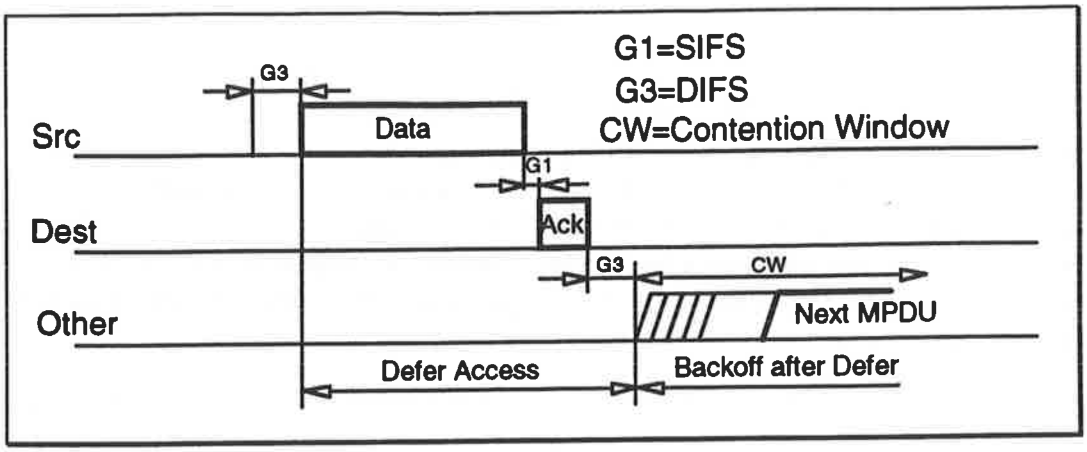

Figure 3-3: Directed Data/Ack MPDU

Note: It is desirable that the PHY has an "End Delimiter" detection means, so that the end of the "Medium Busy" can be detected bit-synchronously, minimizing detection delay..

The following rules need to be applied when transmitters use the Asynchronous Services in the Contention Period:

When transmitting a unicast MPDU using RTS/CTS option:
If medium free (No NAV and no CS) longer then DIFS, then transmit RTS.
Else defer until DIFS gap is detected, and go into backoff.
If CTS is received within T1 after RTS, then transmit the DATA after SIFS.
Else go into Retransmit_Backoff.
If Ack not received within T3 then go into Retransmit_Backoff.

When transmitting a unicast MPDU not using the RTS/CTS option:
If medium free (No NAV and no CS) longer then DIFS, then transmit DATA.
Else defer until DIFS gap is detected, and go into backoff.
If Ack not received within T3 then go into Retransmit_Backoff.

When transmitting a Broadcast/Multicast MPDU:
If medium free (No NAV and no CS) longer then DIFS, then transmit DATA.
Else defer until DIFS gap is detected, and go into backoff.

The following rules need to be applied by receiving stations:

If RTS frame is detected but station is not the destination, Then:
Update the NAV with the Duration information and start a T1 timer.
Else
Return a CTS frame when medium free (no NAV and no CS) after SIFS.

If T1 timer expires, and CS is not active at that time, then clear the NAV.

If CTS frame is detected Then: Update the NAV with the Duration information.

If station is the destination of a unicast DATA frame, Then:
Transmit Ack after SIFS when CRC was correct.

It is required that the transmission of the Ack is done unconditionally. This induces some probability that a pending frame in a neighboring BSA (using the same channel) could be corrupted by the generated Ack. However if no Ack is returned because a CS was detected, then it is guaranteed that the frame would be interpreted as in error, resulting in a retransmission.

3.1.3 Access Backoff Procedure

Stations that want to access the medium, and find the medium busy, will need to defer until the medium is available (after a DIFS gap is detected), and then generate a random backoff period to resolve contention between multiple stations that have been deferring to the same frame occupying the medium.

Backoff= CW * Random() * Slot time

where "Slot Time" = The total propagation delay of :
Transmitter turn-on, medium propagation delay and "Medium Busy Detect" response time.
and CW = A "Contention Window" parameter in slot time intervals.

The Slot Time is dependent on the PHY.

CW is the Contention Window parameter which will have an initial value of CWmin, and which is loaded every time a new frame is put into the transmit queue to be transmitted. The CW should increase exponentially after every retransmission attempt., up to a maximum value CWmax. This is done to improve the stability of the access protocol under high load conditions.

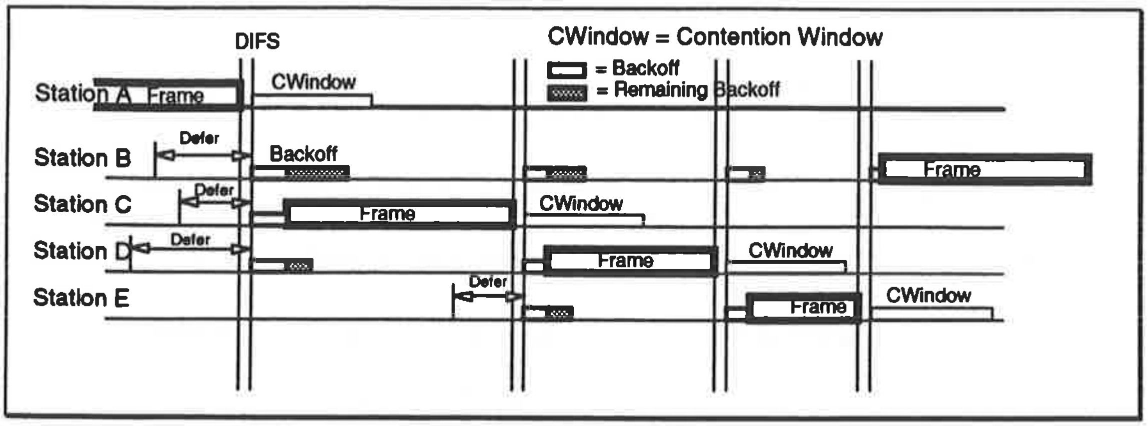

Figure 3-4: Backoff Algorithm

The Backoff delay Timer elapses only when the medium is sensed idle. This means that the Backoff delay timer is frozen while the medium is sensed busy (CS or NAV active), while elapsing is resumed whenever an idle period longer then DIFS is detected.

The effect of this is that when multiple stations are deferring and go into random Backoff, then the station selecting the lowest delay will win the contention. The advantage of this approach is that stations that lost contention will defer again until after the next DIFS period, and will then likely have a shorter Backoff delay than new stations entering the Backoff for the first time.

The Contention Window (CW) is doubled after every retransmission attempt up to a maximum CWmax., to assure stability under high load conditions

3.1.4 DFWMAC Recovery Procedure

Multiple recovery events can be distinguished.

When after an RTS is transmitted, no CTS is being received within a predetermined CTS_Timeout (T1), then a new RTS is generated following the same back-off mechanism as described above, but the ContentionWindow (CW) is doubled up until a CWmax value, at every retransmission attempt.

The same Backoff mechanism is used when no Ack is received within a predetermined ACK_Window (T3) after a unicast Data frame has been transmitted. This process needs to continue until the RE-TRANSMIT_Counter reaches a RE-TRANSMIT_Limit.

Due to several reasons, no CTS may be returned after the RTS transmissions. This can happen due to a collision with another RTS or a DATA frame, or due to interference during the RTS or CTS frame. It can however also be that no CTS is returned because the remote station senses a CS or has an active NAV, indicating a reserved time period.

When no Ack is detected in time, then this can be caused by errors in the DATA or Ack frame due to interference. Therefore different Retransmit limits will be required for retries due to missing CTS responses than for complete retransmissions after no Ack has been received.

Note that the RTS_ RE-TRANSMIT_LIMIT is expected to be set to 8-16, while the no-Ack_Retransmit_Limit is expected to be set to 2-4.

3.1.5 Access Fairness Provisions

A station that has just transmitted a frame and has another frame waiting to transmit, would need to backoff its access also. Otherwise this station would effectively have a much higher access priority, compared with other stations which are waiting to gain access to the medium, and were deferring to the frame that was just transmitted.

This can be accomplished by generating a random backoff delay after every transmission of a DATA frame.

3.1.6 Prioritized Access

Different access priorities can be achieved in various ways. Absolute priority differences that are used for different frame type transmissions are achieved through the use of different Inter Frame Space (IFS) periods between subsequent frames as described earlier. The values for SIFS, PIFS and DIFS are PHY dependent parameters. In addition a high priority DIFS and a low priority DIFS can be specified. The low priority DIFS can be useful for the transmission of certain management frames.

In general different relative priority can be achieved between Asynchronous traffic in the Contention period by using higher DIFS values, or by using different Contention Window sizes. An Access Point can for instance be given some higher access priority by using a lower CWmin parameter then is used in the Stations.

Reference [4] presents some further considerations and possible applications of prioritized access.

3.1.7 CSMA/CA Application

The CSMA/CA+Ack is used for all unicast frame transmissions. This includes transfers between a station and an AP in an Infrastructure environment, and between stations in an Ad-Hoc and Infrastructure network. It is also used between AP's when a wireless Distribution System is being used.

In certain circumstances the DATA frames will be preceded with an RTS and CTS frame exchange that include duration information. A station will use RTS/CTS for directed frames only when the length of the DATA packet is longer then the length indicated by a NoRTS parameter that can be set as a management object on a per station basis. This allows control over the use of this option, such that for instance only inbound long frames use the RTS/CTS option.

CSMA/CA+Ack (with optional RTS/CTS) is also used for all Broadcast frames that are directed towards the AP. So Broadcast/Multicast frames from station to AP, and between AP's will all use CSMA/CA+Ack (with optional RTS/CTS).

For all Broadcast/Multicast frames directed to stations, only CSMA/CA is used, i.e. without acknowledgment and without RTS/CTS. This includes the Broadcast/Multicast traffic from an AP to a Station, or from Station to Station in an Ad-Hoc network. Consequently there will be no MAC level recovery on Broadcast/Multicast frames. Therefore the reliability of this traffic is subject to lost frames due to interference or collisions.

3.1.8 Duplicate Detection

Since MAC-level acknowledgments and retransmissions are incorporated into DFWMAC, there is the possibility that a frame may be received more than once. Such duplicate frames must be filtered out within the receiving MAC.

Duplicate frame filtering is facilitated through the inclusion of an **MPDU ID** field within the individual frames of an MPDU, including the Data and Ack frames. Frames which are part of the same MPDU will have the same ID, and different MPDUs will (with a very high probability) have a different ID.

The MPDU ID is a 16 bit hash of the 2 octet Network ID field, 6 octet source address and a 1 octet sequence number maintained by the transmitting station. The hashing of this information into a smaller field reduces overhead (particularly within Ack frames). A receiving station will reject a frame (which has the "retry" bit set) as a duplicate if it receives one which matches a value of recent MPDU IDs kept within a small cache. There is the small possibility that a frame will be improperly rejected due to such a match;

however, this occurrence would be rare and would simply result in a lost frame similar to a CRC error.

Receiving stations must acknowledge frames even if they are subsequently rejected due to duplicate filtering. Note that this implies that receipt of an Ack is not to be interpreted within the sending MAC as signifying delivery of the frame to the destination MAC user.

Note that in addition to its duplicate detection function, the MPDU ID construct is used to tie together the different components of an RTS/CTS exchange as well.

3.1.9 Fast Response Possibility

Note that instead of an Ack frame, it is also possible to directly transmit the response frame back to the transmitter of the received frame. This would allow a class of fast implementations, which could for instance directly respond to a Poll frame with the requested Data frame itself, which in turn should be Acknowledged by an Ack frame. Another example is in the Contention Free (CF) period, where stations respond to a Poll bit in frames coming from the AP.

3.2 Point Coordination Function

DFWMAC optionally supports a Point Coordination Function (PCF), which can provide Contention Free services. The use of this PCF is restricted, because it can only be used in certain environments.

The basic restriction is that a PCF can not overlap with another PCF on the same channel, so sufficient isolation between multiple PCF's is needed. This is because contention between multiple overlapping PCF's can not be resolved by the protocol. This will limit the useability of the PCF to either operation in a single BSS of an ESS, or to multiple channel environments that can assure sufficient isolation between neighboring PCF's.

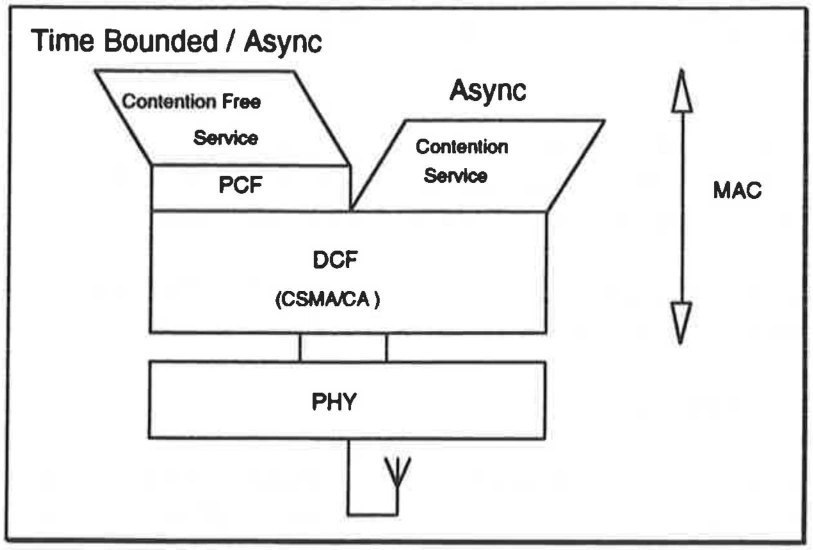

Figure 3-5: MAC Service Model

The PCF allows Contention Free services, which can be utilized for both Time Bounded and Asynchronous Services. The Contention Free services are only supported in infrastructure mode, but it can support Contention Free direct Station-to-Station services in this mode. Stations and AP's have the option to implement the PCF. They can only be used when both a station and an AP has that option implemented.

As shown in Figure 3-5, the optional PCF is build on top of the CSMA/CA based DCF, by utilizing the access priority provisions provided by this scheme, as explained in Section 3.1of this document.

The PCF can support Contention Free Time Bounded and Asynchronous services.

3.2.1 SuperFrame Structure

The PCF uses a Superframe concept as shown in Figure 3-6. Within a given SuperFrame period, the PCF will be active in the Contention Free Period, while the DCF is used in the Contention Period.

The Contention Free Period can be variable in length on a per SuperFrame basis, without any additional overhead. At the beginning of the SuperFrame, the PCF will gain control over the medium provided that the medium is sensed free. If the medium is sensed busy, then the PCF will defer until the end of the frame, and will gain control over the medium as soon as it senses the medium available for longer then the PCF period. This will cause stretching of the SuperFrame, causing a variable start of the Contention Free period.

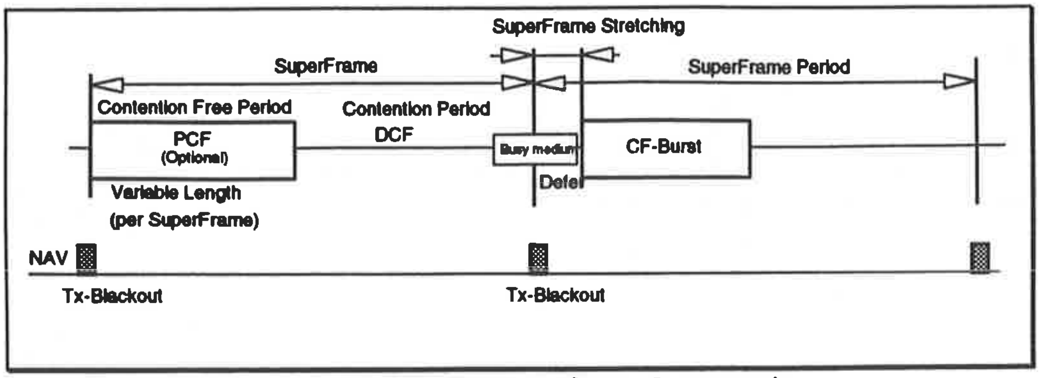

Figure 3-6: PCF Superframe Construction

The Asynchronous traffic that uses the DCF will automatically defer until after the Contention Free Period, because the PCF uses the PCF priority level of the CSMA/CA access protocol, which causes a burst of traffic with interframe gaps that are smaller then the minimum DIFS period needed by the CSMA/CA protocol.

To prevent contention at the start of the SuperFrame period between the Contention traffic and the PCF, a Tx-Blackout function can be used. This can be seen as a preset of the NAV of the first slot at the start of every SuperFrame.

The length of a SuperFrame could be a manageable parameter, that may depend on the services provided, and perhaps the type of PHY used. In a Frequency Hopping PHY it is required that the Hop Dwell time is an integer multiple of the SuperFrame Period.

3.2.2 PCF protocol

The PCF protocol is based on a polling scheme. Stations need to request the Contention Free Service, which when granted can then be used until the service is closed.
If the service is granted by the AP, then the PCF protocol will place that station in the Polling list, such that it can participate in the service.

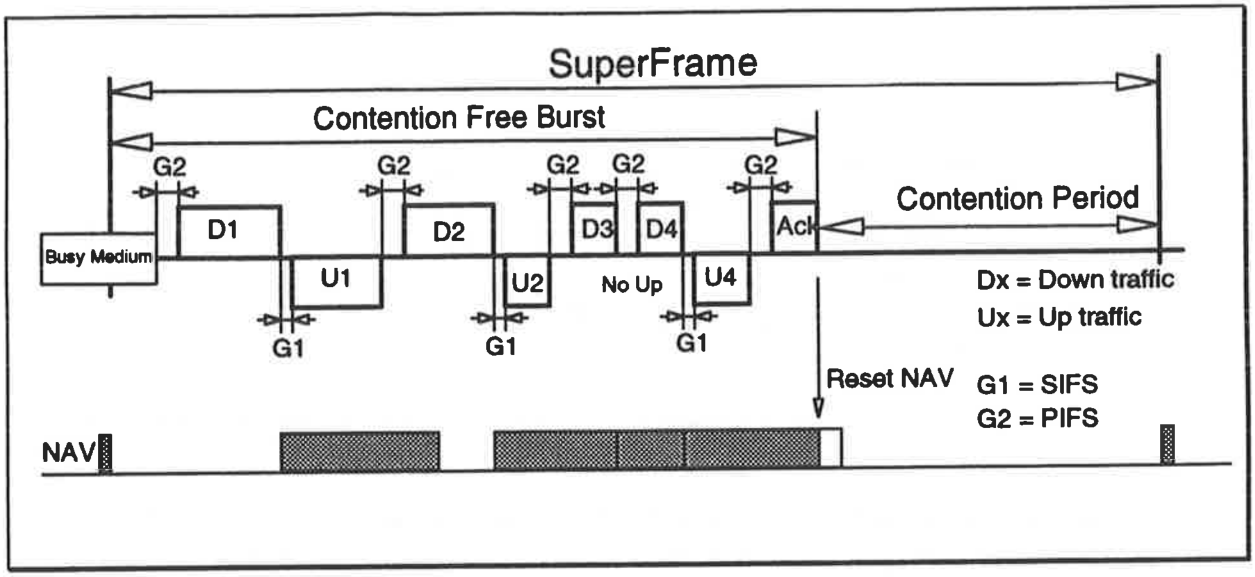

Figure 3-7: PCF Protocol Mechanism

AP's will send the buffered (CF-Down) traffic for Contention Free stations at a specific PCF priority (PIFS gap). A Polling bit in the Header of these frames will poll the stations to send their (CF-Up) data if any. Stations will react on the poll immediately when they have a frame queued up for transmission in the Contention Free Period, by sending this frame after the SIFS time gap. This results in a Burst of Contention Free traffic, the CF-Burst.

For services that require MAC level acknowledgment, the Acknowledgment is done through a bit in the Header of the responding or subsequent frames. This is the Ack_Previous_Frame or APF bit. So the U1 frame in above figure contains the APF bit to acknowledge the previous D1 frame. Also the D2 frame will contain the APF bit to acknowledge the superseding U1 frame.

The duration of the frames can be variable, only bounded by a maximum frame length negotiated during the Contention Free Service request. If a station does not react within the PIFS delay time, then the PCF in the AP will resume control and transmit the next queued frame. Note that a station will not respond when:

> The station has no CF-Up traffic to send, **AND** no Ack is required to be returned for the preceding CF-Down frame. No Ack is for instance needed when the CF-Down frame had a Data field length of 0.

In the Header of all the Contention Free frames coming from the AP (CF-Down) an indication is given to update the NAV in every station deferring for the CF-Burst. This will prevent that stations that do not hear a (hidden) responding station, will inadvertently start their transmission during the CF-Burst.

Note that the assumption is that all stations will sufficiently hear the AP, so that the NAV information that is received by every station can then protect the CF-Up frames, which may not be heard by all the stations in the BSS.

The PCF will generate a frame at the end of the CF-Burst that has the following functions:

- Acknowledge the last CF-Up frame if needed through the APF bit.
- Reset the NAV in all stations that are listening.

This can be done by for instance an Ack or a Beacon frame. It is important to reset the NAV, because of the variable size of the Station response frame. The NAV is set to a maximum expected response frame length. A Reset_NAV function is needed to enable all Asynchronous stations to use the contention period immediately following the end of the CF-Burst.

3.2.3 PCF Station-to-Station support

The PCF does also support Station-to-Station transfers as is shown in the figure below. A separate Poll frame that is addressed to a station will trigger that station to generate its frame after the SIFS. This can be a frame that addresses another station instead of the AP. If needed, then the addressed receiving station can generate an Ack to acknowledge the proper reception. The PCF will resume control over the medium and send the next CF-Down frame after a G2 gap.

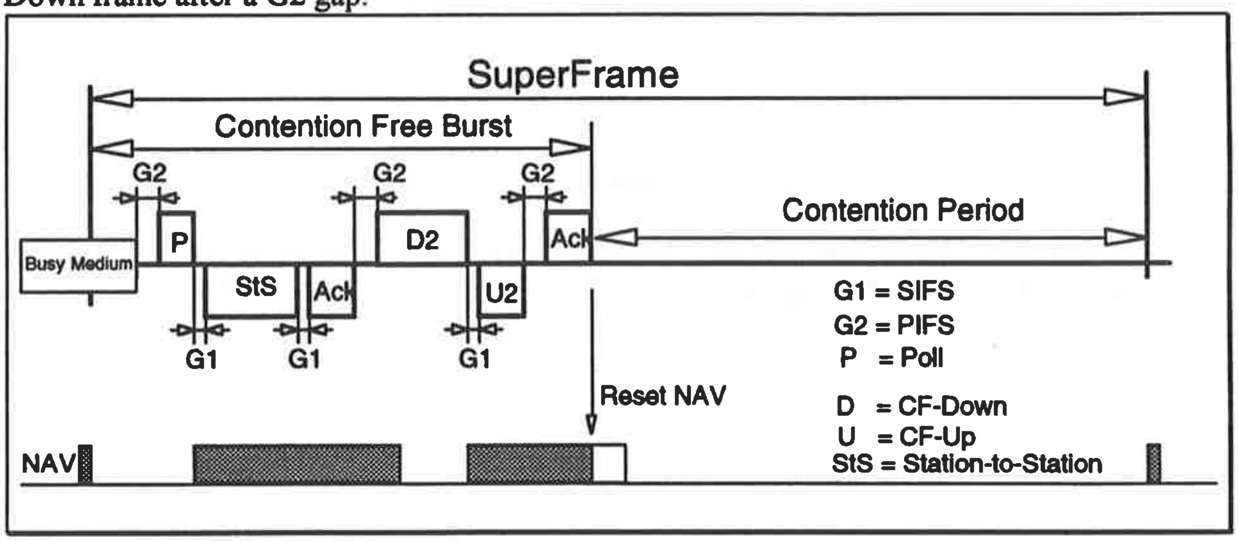

Figure 3-8: Station-to-Station PCF Protocol

3.2.4 Contention Free Service Types

The following two different PCF Service Types are provided, that can have different service levels:

- Time Bounded Services (TBS = Time Bounded Service)
- Asynchronous Services (ACFS = Async Contention Free Service)

They will both use the same PCF, but use different reservation schemes, different addressing mechanisms, and different frame formats. The PCF will use different allocation mechanisms for these types of contention Free Services.

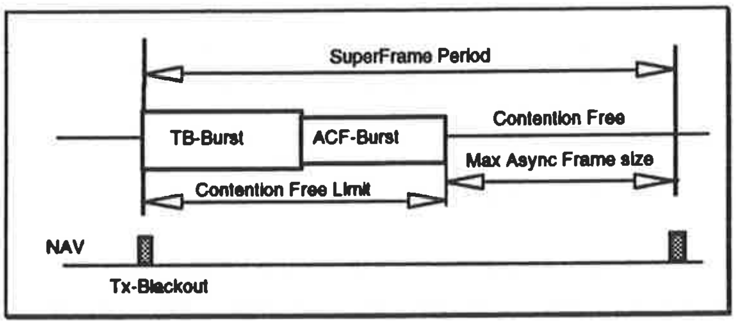

Figure 3-9: PCF Contention-Free Limits

The Contention Free period needs to be limited to allow coexistence between Contention and Contention Free traffic. The absolute maximum time that is allowed to be allocated to these services in a SuperFrame is such that at least one maximum size Asynchronous MPDU can still fit into the SuperFrame period. So:

Contention Free Limit = SuperFrame Period - Max. Async MPDU size

This will allow that at least one Asynchronous MPDU can be transmitted in the contention free period, without affecting the timing constraints of the Time Bounded Service. Note that the start of the TB-Burst can jitter due to the SuperFrame stretching that can occur when the PCF must defer for pending DCF traffic.

The PCF needs to apply a limit for the Time Bounded traffic, such that enough time is available for the Contention traffic, while also providing for a certain amount of Asynchronous Contention Free capacity.

Section 6 will discuss the two different types of Contention Free services that are provided by the optional PCF.

3.2.5 Scan and Power Management Provisions

Provisions are included to allow a station sufficient time to "Scan for a better AP". This is done by randomizing the sequence in which the different stations are serviced within a SFP. This will allow all stations on average the same time per SuperFrame to perform the scanning process until a new SFP is started.

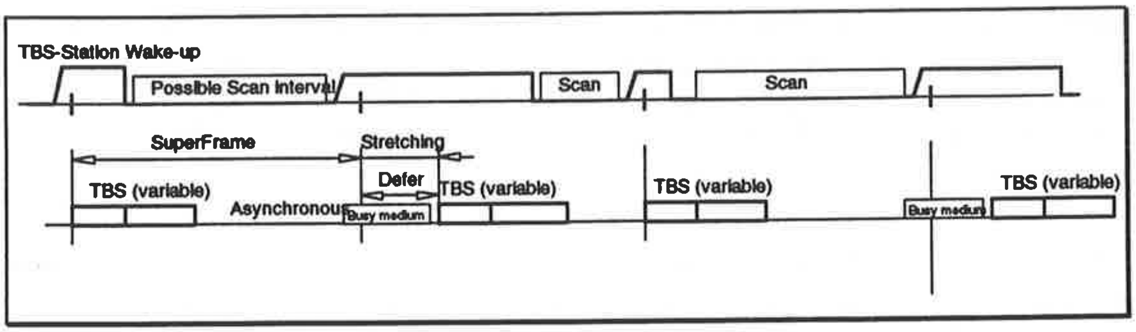

Figure 3-10: Scan and Power Management Provisions

The same provisions will also provide all stations using the PCF equal opportunity to save power by turning off their transceiver after the station has been polled, until the end of the SFP.

4. Synchronization

It is beneficial for a variety of reasons to synchronize all nodes within a single BSS to a common clock. As we will describe in this section, synchronization does not require a synchronous protocol in which the time domain is partitioned into fixed frames and certain transmissions must occur at designated times. Such an approach would be in conflict with a CSMA access method. The method proposed here to achieve synchronization is compatible with the CSMA approach and does not violate the basic CSMA rule which causes stations to defer their transmissions upon perceiving a busy medium.

4.1 Basic Approach

All stations maintain a local synchronization timer. The synchronization timer is used for power management, superframe timing, Frequency Hop synchronization, and other management functions that require periodic activity of the station. A Timing Synchronization Function (TSF) keeps the timers for all stations in the same BSS synchronized.

4.1.1 TSF for Infrastructure Networks

In an infrastructure network, the AP is the timing master and performs the Timing Synchronization Function. To synchronize the other stations in a BSS, the AP sends out periodic frames called Beacons that contain a copy of its timer. Receiving stations always accept the timing information in Beacons sent from the AP servicing their BSS. If their local timer is different from the timestamp in the received Beacon, they set their local timer to the received timestamp value.

4.1.2 TSF for Ad Hoc Networks

The Timing Synchronization Function in an ad hoc network is a distributed algorithm that is performed by all of the members of the BSS. This is important to maintain the spontaneous nature of ad hoc networks. Stations may join and leave the network in any order without re configuring the network. Since there is no master node, all stations in the BSS must share the responsibility for sending out periodic Beacons. Stations in the same BSS adjust their local timers towards those that they hear in Beacon frame sent by other stations in the BSS. The timers for all synchronized stations in a BSS will converge on the same value over a short period of time.

4.2 Timing Synchronization Function Applications

With all stations within the BSS synchronized, network events can be scheduled to occur at specific times. This may involve the transmission of a particular type of frame by a certain station (sending a Beacon or starting the Time Bounded burst), or the transition of the entire BSS into a new state (moving to the next frequency in a FHSS). The following are example applications which can be accommodated within a synchronized BSS:

- Power management, allowing nodes to power off their receivers until the expected transmission of the next Beacon.

- PHY management, determining the hop timing in a Frequency Hopping PHY.
- Time Bounded traffic support, defining and maintaining superframe timing.

It is important to point out that even when the TSF is used to schedule the transmission of a specific frame, it is not necessary that the frame be transmitted at that exact instant. Instead, the frame's transmission can be deferred as usual according to the basic CSMA algorithm. The scheduled transmission time therefore should be viewed as a earliest transmission time.

Later sections of this document will describe how the TSF is used to implement power management and define the superframe for Time Bounded Services. First we describe how the basic synchronization within a BSS is achieved even if stations are operating in a low power mode with their receivers turned of for long periods of time.

4.3 Maintaining Synchronization

Each station has its own internal free-running timer with modulus TSFTIMERMOD. Stations expect to receive Beacons at a nominal rate. The interval between Beacons is defined by the BEACON_INTERVAL parameter of the station. Whenever a Beacon is received, stations may adjust their local timer to match the TSF timestamp contained in the Beacon.

4.3.1 Beacon Generation in Infrastructure Networks

The access point defines the timing for the entire BSS. Maintaining synchronization is simple. The AP sends out Beacons at the desired Beacon interval as defined by its internal clock. Beacons are not sent relative to the transmission of the last Beacon. Though the transmission of a Beacon may be delayed because of CSMA deferrals, subsequent Beacons will be scheduled at the normal beacon interval. This is shown in figure 4.1.

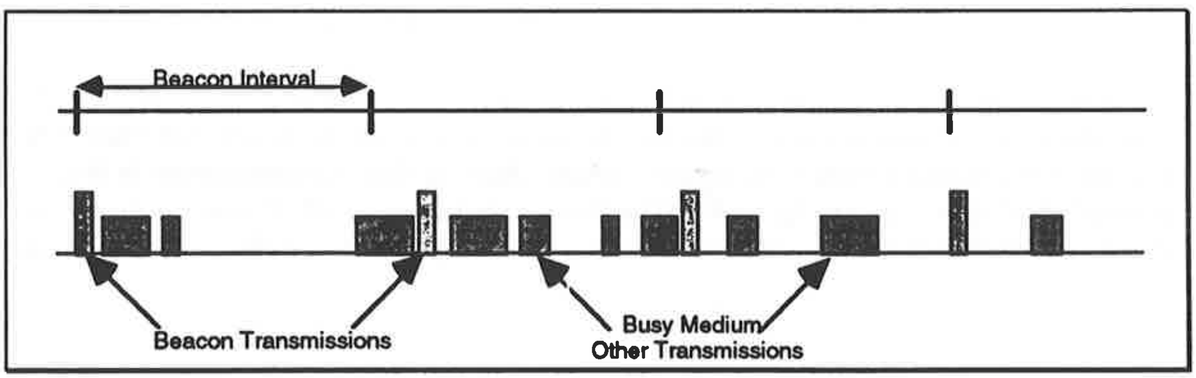

Figure 4.1 — Beacon transmission on a busy network.

4.3.2 Beacon Generation in Ad Hoc Networks

Beacon generation in an ad hoc network is distributed. All members of the BSS participate in Beacon generation. Each station maintains its own TSF timer which is used for BEACON_INTERVAL timing. Each time the BEACON_INTERVAL time elapses, the station will wait a random delay then send a Beacon. If a Beacon is received from another station during this process, the transmission is canceled.

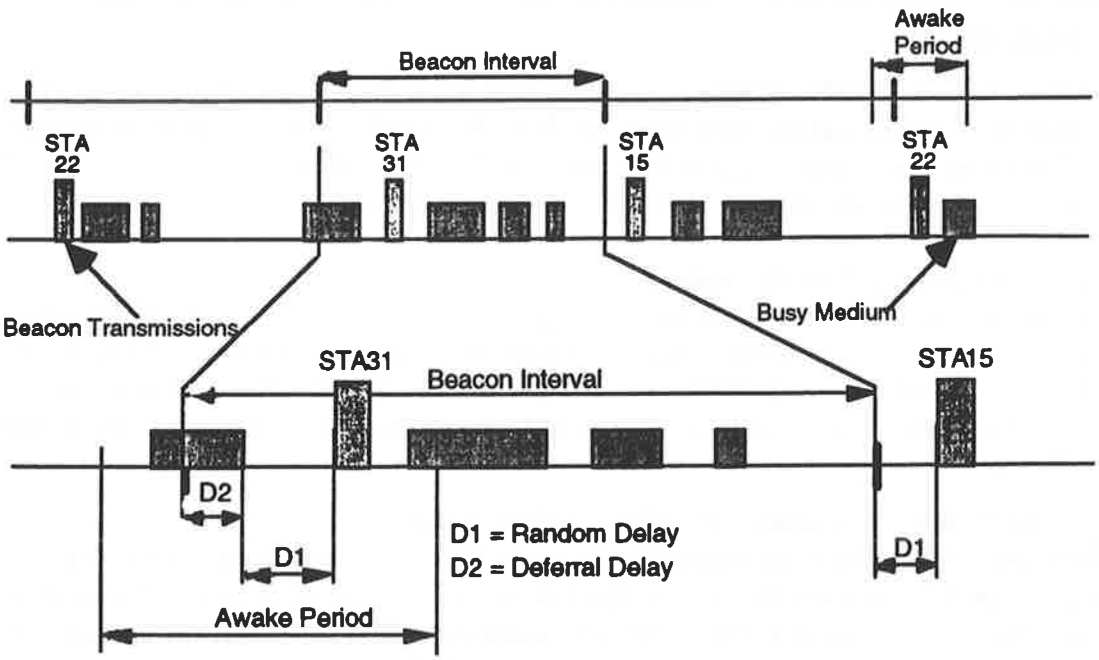

Figure 4.2 — Beacon transmission in an Ad Hoc network.

The randomized Beacon generation delay uses the same technique as the Access Backoff Procedure that stations use when they attempt to transmit and find the network busy. The random delay happens even if the network is idle when the Beacon transmission is first attempted. This is shown in figure 4.2. The Beacon transmission will always occur during the Awake Period of stations that are operating in a low power mode. This is described in more detail in section 5.

4.3.3 Synchronization Beacon Content

A synchronization Beacon contains a message header and the following required elements:

- Timestamp — a 32 bit field containing the TSF timer MOD 2^{31} (in microseconds) and a Boolean flag which is 1 if the station sending the Beacon is synchronized. An AP is always synchronized. Ad hoc stations may not be.
- Beacon interval — a 24 bit field containing the time between Beacons.

Beacons are sent to the broadcast destination address and contain the NID of the sending station. A Beacon from an AP in an infrastructure network may contain other optional elements that are used for Power Management:

- TIM — traffic indication map.
- DTIM period — number of TIMs between each DTIM.
- Broadcast Element

A Beacon from a station in an ad hoc network will contain:

- Weight — a 16 bit value. The Weight of the sending station.

> The Weight variable is used to assist in adjusting synchronization timers in ad hoc networks. Stations may hear Beacons from other stations in the same BSS that have different timestamp values. When this occurs the receiving station may adjust its timer towards that contained in the Beacon. The Weight variable determines how much to adjust the timers. The longer a station has been in a BSS and the more stations it has heard the more Weight it should have. A new station entering the network should have less influence on the receivers' synchronization timer. A possible formula for determining the Weight is:
> Weight = #Stations_heard * #Beacons_heard.
>
> #Stations_heard is the number of different stations within the same BSS from which we received Beacons.
>
> #Beacons_heard is the number of Beacons we have heard during the last 10 Beacon intervals.

4.3.4 Synchronization Timer Accuracy

The TSF timer has a resolution of microseconds. It is important to keep all of the timers within a BSS synchronized. Beacon transmissions may have to defer to network traffic. To ensure that the timestamp value in the Beacon accurately represents the internal timer, the timestamp field is not filled in until after the CSMA deferral. The start of the MAC frame is used as the timing reference. The timestamp value in the Beacon frame is the value of the free-running synchronization timer at the instant that the Start Frame Delimiter (SFD) is transmitted.

Stations receiving the Beacon, save the value of their local timer at the instant the SFD is received. If the Beacon is from their BSS and has a valid CRC, they will accept the Beacon. Stations can then accurately update their synchronization timer by using the difference between their local timer when the SFD for the Beacon occurred and the timestamp field received in the Beacon. Beacons are never forwarded beyond a BSS, so the delay between the Beacon SFD transmission and the reception of the Beacon SFD is simply the propagation delay and a possible transceiver processing delay. It is possible therefore to obtain very precise synchronization of the TSF timers for all stations within a BSS. It is not necessary that a station properly receive every Beacon to maintain accurate timing.

4.4 Acquiring Synchronization, Scanning

Stations wishing to join a BSS must find the other stations belonging to that BSS by listening on the appropriate channel and synchronizing their TSF timers with the rest of the BSS. This is usually done by scanning all of the possible frequencies for a period of time listening for messages from the AP or other stations in the BSS.

4.4.1 Passive Scanning

Since Beacons are transmitted periodically, it is possible for a new station attempting to join a network to simply listen on each possible frequency until a Beacon is heard. If a Beacon is heard, it contains all of the information necessary to synchronize with the rest of the BSS. Passive scanning is possible if the BEACON_INTERVAL is relatively short or if there are just a few different channels supported by the PHY.

4.4.2 Active Scanning

To speed up the scanning process, there is an additional mechanism that enables stations to quickly learn about their environment. One of the MAC Management messages is a Probe Request (simply called a Probe). If an AP hears a Probe it will send a Probe Response. In an Ad Hoc network the responsibility for sending the Probe Response is handled in a manner similar to Beacon generation.

A Probe is simply a message header with the appropriate message type. Probes are usually sent to the broadcast address and have an NID with a specific ESSID and either a broadcast or specific BSSID.

Probe Responses are always sent to the source of the Probe as directed messages and contain an NID matching the Probe. A Probe Response also contains the following elements:

- Timestamp — a 32 bit field containing the TSF timer MOD 2^{31} (in microseconds) and a Boolean flag which is 1 if the station sending the Beacon is synchronized. An AP is always synchronized. Ad hoc stations may not be.
- Source NID — a 24 bit field containing the NID of the sending station. The NID in the header always matches the Probe, so it may be a broadcast BSSID.
- Weight — a 16 bit value. The Weight of the sending station.

4.4.3 Sending a Probe Response

There may be multiple stations responding to a Probe. To avoid an instant collision, stations responding to Probes must follow these rules. Always send a probe Response as a directed message. Always defer and do a Backoff before sending a Probe Response even on the first transmission attempt. This relies on the basic deferral and backoff algorithms already built in to the protocol to minimize the probability of a collision and retransmit the Probe Response if a collision does occur.

4.4.4 Scanning for Infrastructure Networks

APs are always in Sync and always respond to Probes. It is possible that more than one AP is active on the same channel. In that case, it is desirable for the station sending the Probe to hear responses from all APs that are on that channel. The scanning station may be trying to determine which AP is bètter, and the quality of the Probe Responses is a simple way to compare one against the other.

Figure 4.3 shows two APs on the same channel responding to a Probe. AP 2 picks a shorter Backoff time and sends a Probe Response first. The scanning station ACKs the Probe Response. AP 1 had deferred to AP2's transmission, and it's Backoff timer elapses shortly after DIFS gap following the ACK. AP 2 then sends the Probe Response which is also ACKed by the probing station.

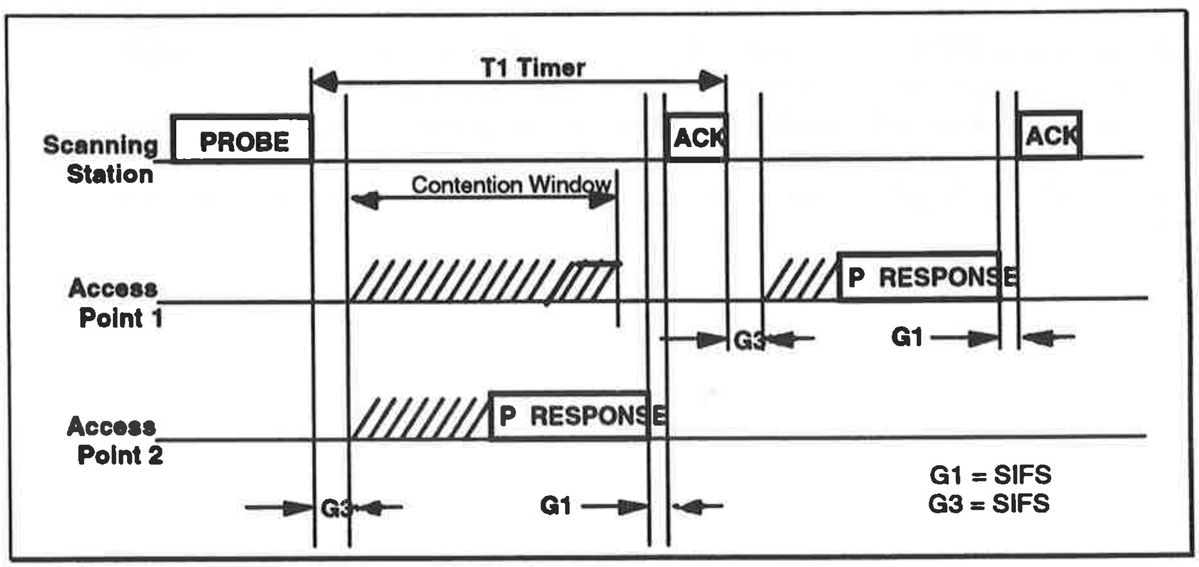

Figure 4.3 — Multiple APs responding to a Probe.

In the example above, there was no collision. Both Probe Responses were sent without error. If a collision had occurred, the probing station would not send an ACK. The APs would each schedule a retransmission with a Contention Window that is twice as large which would resolve potential Probe Response collisions. Though the scanning station will not be able to interpret the collision, it will stay on the channel longer because the presence of energy on the medium indicates this channel is a "likely candidate".

In order for this technique to work effectively, a scanning station should take the following actions to find any AP belonging to a certain ESS :

1. CSMA defer
2. Send Probe with Broadcast Destination, Specific ESSID, and Broadcast BSSID
3. Start T1 timer
4. If receive energy within T1 then start T2
   Else If hear nothing within T1
   then Scan next channel.
5. When T2 expires, process all received Probe Responses.
6. Scan next channel.

Whenever a station switches to a new channel, the information in its NAV is invalid until it hears at least one MPDU on the new channel. Stations doing fast scanning simply clear the NAV each time they switch to a new channel. Therefore the Probe is sent as soon as Carrier Sense indicates that the network is idle. This significantly reduces the time required to check an empty channel. T1 defines the length of time a scanning station will have to wait on an empty channel before concluding that there will be no response. T2 is longer, and is the length of time a scanning station will wait on a channel when it has detected energy.

As an example of fast scanning, let 's examine the time required to determine that no stations are present on a network. Assume a 1 Mbps frequency hopping system with slow channel switching time. The scanning station will send a Probe and wait T1 time on each channel before concluding that it is empty.

Scan Time per Empty Channel = Probe Time + T1 + channel switch time
Probe Time = 300 μseconds (30 byte packet plus preamble)
T1 = Contention Window + gap
Contention Window = 32 * Slot Time = 32 * 40 μsecs = 1280 μsecs
gap = 100 μsecs
Channel Switch Time = 300 μseconds
Scan Time per Empty Channel = 300 + 1280 + 100 + 300 = 1980 μseconds

This hypothetical frequency hopping system has 75 channels. It is possible to scan all channels in the system in less than 150 milliseconds. To account for the random hopping nature of FH PHYs, a scanner might check all channels three times. This would take a total of 450 milliseconds.

Using DFWMAC's fast scanning, a station attempting to find a network in a system with these PHY characteristics can conclude that there is no other station within range in less than a half of a second.

4.4.5 Scanning for an Existing Ad Hoc Network

A new station attempting to join an existing ad hoc network will scan for the network in a similar manner. The scanner uses the following procedure:

1. CSMA defer
2. Send Probe with Broadcast Destination, Specific ESSID, and Specific BSSID
3. Start T1 timer
4. If receive energy within T1
   then start T2
   Else If hear nothing within T1
   then Scan next channel.
5. When T2 expires, process all received Probe Responses.
6. Scan next channel.

The only difference is the second step. The scanning station wishes to join a particular ad hoc network, so it specifies the BSSID and ESSID in the Probe message.

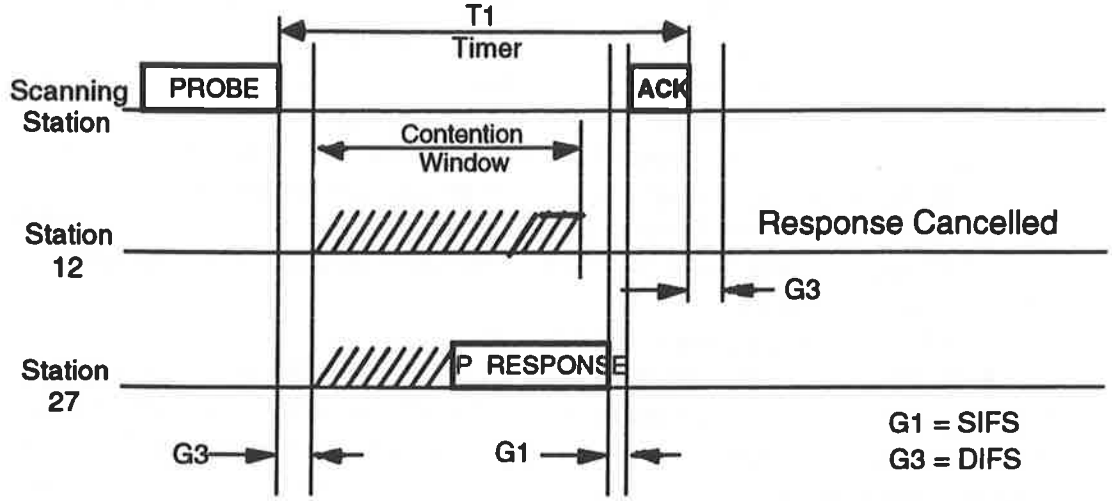

Figure 4.4 — Probe Response Contention in an Ad Hoc network.

A Probe may elicit multiple responses as in the infrastructure case. The contention is resolved in the same manner. All stations responding to the Probe behave as if they sensed the media busy just before transmitting the Probe Response. They do a random backoff before attempting to send. Probe Responses are sent as directed messages. If a collision does occur, the sending stations will not receive an ACK and will retransmit the Probe Response again with a larger Contention Window.

It is not necessary for a scanning station to hear every member of an ad hoc BSS. Whenever a responding station hears a Probe Response from another station followed by an ACK, it will cancel its own Probe Response transmission. This ensures that a scanning station will get at least one Probe Response. Figure 4.4 describes this behavior. The scanning process stops when a Probe Response is heard with the desired NID. The scanning station can join that network by adopting the same TSF timer and NID.

It is possible that the ad hoc network is operating in a low power mode. Stations belonging to a low power BSS will have their receivers on for short awake periods around the expected transmission of Beacons. In that case, scanning could take much longer especially for a Frequency Hopping system. To improve this situation, ad hoc stations sacrifice some of their power. Each time a station joins the network it selects a small set of the possible channels. Whenever the current hop is one of those channels, that station stays active for the entire dwell time so that it can respond to Probe requests. This increases the probability that a new station will send a Probe when someone from the ad hoc network is active. For a single channel PHY, the same effect is possible if stations periodically stay awake for an entire BEACON_INTERVAL.

4.4.6 Initializing an Ad Hoc Network

The procedure for starting an ad hoc network is as follows:

1. Scan (BSSID = specific) for 5 seconds
2. If hear something then JoinNet
3. else StartNet
4. Stay active (don't go into low power mode) for T3 seconds

JoinNet is the process described above as scanning for an ad hoc network. StartNet causes the station to establish the initial timing reference for the ad hoc network. The station will set its SYNC flag to 1 and begins sending Beacons at the expected intervals. Stations following the same process slightly later than the first station will find the first station during the long Scan in step 1. There is a slight probability that two stations can execute the process at the same time and start two independent clusters of the same network. Clusters can also form when stations out of PHY range from each other start the same network and then some members of the different clusters move within range of each other.

Clusters are simply groups of stations using the same NID that can not communicate with each other because they are unsynchronized. They may be unsynchronized in terms of their ad hoc superframe timing so that their awake periods never coincide. Clusters may also be operating on different channels in a multi channel PHY.

4.4.7 Other Applications of Scanning

So far, we have described scanning as a method of finding and synchronizing with a network when a station initializes itself. Scanning can also be used to solve the clustering problem for ad hoc networks and assist stations in an infrastructure network transition from one BSS to another. In both cases, stations should occasionally scan to find other BSS or other clusters of the same BSS. In a network of stations operating in a low power mode, stations can scan in between awake intervals. This is described in figure 3-10. in a continuously active network, stations may scan when they believe they will be idle.

When scanning stations find another BSS or cluster they can save the difference between the TSF timer on the scanned network and the TSF timer on the home network as a description for how to find the other network. For networks that are already synchronized the difference between their TSF timers should be fairly constant over a long period of time.

4.5 Adjusting Station Timers

Stations always adopt the timer in a Beacon or Probe Response coming from the AP in their BSS. In an ad hoc network, the rules are not so simple. A station must determine whether to adopt the timer from an incoming Beacon or Probe Response; or to adjust its own timer in the direction of the timestamp received. The SYNC flag and Weight variable are used for this purpose.

When a station first enters a network, it is unsynchronized (its SYNC flag is 0). After a station has heard a number of Beacons from other stations that are synchronized and the timestamp in the Beacons matches the station's timer; it changes to the synchronized state. If a station adopts the timer of an AP or a synchronized station, it changes its state to synchronized.

The timer adjustment rules are as follows:

```
If the receiving station is unsynchronized:
    if receive Beacon or Probe Response with SYNC = 1 then adopt timestamp
    if receive Probe Response with SYNC = 0 AdjustTimer
If the receiving station is synchronized:
    if receive Beacon or Probe Response with SYNC = 1 then
        if the difference between the timers is above a threshold Coalesce
            else AdjustTimer
    if receive Probe Response with SYNC = 0 ignore
```

AdjustTimer uses the Weight variable in the Beacon or Probe Response to determine how much to adjust the station timer. If an incoming Beacon contains a Weight that is greater than the station Weight, the station timer is adjusted by a large amount towards the timestamp contained in the Beacon. By contrast if he station weight is greater, then the station timer is adjusted by a small amount.

The Coalesce operation needs further study. The basic concept is that a station finding a different cluster will initiate some action that will cause the two clusters to coalesce. This involves determining which cluster is the desired cluster and then sending a broadcast Management message (a Beacon?) on the undesirable cluster to inform all members of that cluster to synchronize with the other cluster.

4.6 Timing Synchronization for Frequency Hopping PHYs

The TSF described here provides a mechanism for stations in a frequency hopping system to synchronize their transitions from one channel to another (their hops).

Every station has a table of all of the hopping sequences that are used in the system. All of the stations in a BSS use the same hopping sequence. There is a function that maps from BSSID to a particular entry in the hopping sequence table. This is used to select the hopping sequence for a BSS. Every frame contains the BSSID, so a new station attempting to join a BSS can determine the hopping sequence as soon as it hears any frame. Since it knows the frequency on which it is receiving, the station will be able to determine which hop in the sequence is the current hop and the frequency to use for the next hop. The only missing information is timing — when to tune to the next hop.

Stations can use their TSFTIMER to time the DWELL_INTERVAL. The DWELL_INTERVAL is the length of time that stations will stay on each frequency in their hopping sequence. Once stations are synchronized, they have the same TSFTIMER value. The TSFTIMER is a 31 bit value measured in microseconds. In a frequency hopping system, the maximum value of the TSFTIMER is the number of microseconds it takes to complete one hopping sequence — DWELL_INTERVAL * #Hops.

Stations in the BSS tune to the next frequency in the hopping sequence whenever:

TSFTIMER MOD DWELL_INTERVAL = 0

The time remaining in the current hop is:

DWELL_INTERVAL - (TSFTIMER MOD DWELL_INTERVAL).

When a station is scanning for other networks in a frequency hopping system it will learn all of the essential information about the other network by simply hearing a Beacon or Probe Response. The MAC header contains the BSSID which maps into a hopping sequence and the Timestamp element describes which hop in the sequence and the time remaining in the hop. If the scanning station records the difference between the Timestamp and its local TSFTIMER and the BSSID; it will be able to instantly tune to the other BSS. This is essential to support fast BSS transitions in a frequency hopping system.

This technique for frequency hop synchronization does not require a PHY specific field in the header of every frame. It is not necessary to receive a Beacon in every hop and there is no explicit message that causes stations to switch to the next hop.

5. Power Management

5.1 Power Management in Infrastructure Mode

A wireless station will consume power through both transmission and reception. Since transmitters are powered on only when an actual transmission is occurring, no further reduction in power consumption is possible for transmission. However, receivers are typically powered on even when no actual reception is occurring, to ensure that if a frame does appear it will be properly handled. This presents an opportunity for power savings if the receiver can be powered off except when actually needed.

The mechanism described in this section allows stations to operate with receivers powered off for much of the time while still ensuring that frames will be received. In general there is an inverse relationship between the power consumed and the performance achieved, and it will be seen that this scheme allows stations to operate at various points on the "power/performance" tradeoff curve, independent of other stations. Also, the scheme does not rely upon higher-level software making the receiver on/off decision, and hence is application-independent.

5.1.1 Overview

The basic approach is for stations to inform the AP if they wish to enter a power-conserving mode. In this case, the AP will not arbitrarily transmit frames to the station, but will buffer frames and will only transmit them at designated times.

The stations which currently have buffered frames within the AP are identified in a **Traffic Indication Map**, or **TIM**, which is periodically generated by the AP as an element within a beacon. A station may determine that a frame is buffered for it by listening to a TIM, and can then take action to ensure that the frame is received. The beacon also includes a timestamp which is used for synchronization, as described in Section 3 above.

Beacons (with TIMs) are periodically generated by the AP, separated by a time interval called the **Beacon_Interval** (which is a BSS parameter known by all stations within the BSS upon association with the AP). Stations can individually determine how frequently to listen for beacons, given by a per-station **Listen_Interval** parameter, the value of which may depend upon the station's desired power/performance target. Since beacons are generated at regular intervals (subject to CSMA deferral as described in Section 3), stations may synchronize their Listen_Interval with the Beacon_Interval so as to ensure that the Beacon is likely to arrive immediately after the receiver is powered on.

To accommodate the wide range of anticipated station requirements as regards the tradeoff of power consumption versus performance, the protocol allows for two ways in which a station may obtain its buffered frames. Stations operating in the **Power Save Polling** mode (PSP) will transmit a short Poll frame to the AP, which will respond with the corresponding buffered frame. Such stations may choose to listen to each TIM or may choose to listen only to occasional TIMs, depending upon their desired level of power consumption and performance. Other stations may operate in the **Power Save Non-Polling** mode (PSNP), in which case they listen for certain specific TIMs (called **Delivery TIMs**, or **DTIMs**) after which the AP will deliver their buffered frames without waiting for a poll. Broadcast or multicast frames are delivered to all stations following the DTIM transmissions.

It is anticipated that different networks may set the Beacon_Interval to different values depending upon the nature of the applications being supported. It is likely that most applications would be well served with Beacons (including TIMs) generated every 20 to 50 milliseconds, and with DTIM-containing Beacons every 50 to 200 milliseconds.

5.1.2 Station Power Save Modes

A station transceiver can be in three different states:

Transmit:	Transmitter is turned on.
Awake:	Receiver is fully powered.
Doze:	Transceiver is not able to transmit or receive and consumes very low power. Some circuitry (like timers) may still be active.

Stations which are not concerned with power conservation will likely cycle between Transmit and Awake states. A non-transmitting station which is trying to conserve power, however, will go to the Awake state only if there is a strong likelihood that a frame is about to be received, and otherwise it will reside in the Doze state.

Given these transceiver states, various **power-save modes** for a station are defined. The modes differ in the manner that the station cycles through the transceiver states, and in the actions that the station takes in order to receive a frame.

The first class of stations are typically those which are not battery powered, which can always leave their receiver on and are said to operate in a *Continuous-Active-Mode* (CAM). Such a station will likely achieve the highest level of performance. Battery powered stations can choose to operate in such an activated state temporarily (this is the *Temporary-Active-Mode, or* TAM mode) to allow higher performance operation. Ordinarily a station which wishes to conserve power will operate in either the *Power-*

Save-Nonpolling or *Power-Save-Polling* mode. These modes are summarized in the table below.

Continuous-Active-Mode or CAM	Station may receive frames at any time, no AP buffering
Temporary-Active-Mode or TAM	Similar to CAM mode but temporary, no AP buffering
Power-Save-Non-Polling or PSNP	Station listens to all DTIMs and keeps its receiver on if DTIM indicates a frame is buffered. AP transmits buffered frame without waiting for a poll after the DTIM is transmitted.
Power-Save-Polling or PSP	Station listens to selected TIMs (based upon its Listen_Interval) and polls AP if the TIM indicates a frame is buffered (except for broadcasts). AP only transmits buffered frame in response to poll.

Note that the PSP mode is parametrized by the station's Listen_Interval, which determines how frequently the station listens for TIMs. This mode thus covers a wide range of operational requirements for different stations. For example, devices with extreme power conservation requirements (such as PDAs) can be supported with the PSP mode using a relatively long Listen_Interval parameter. Note that this may impact the higher level protocols which can be supported, and may also impact whether or not Broadcast/Multicast services are available to these stations.

Stations may transition between modes if desired. Transitions between any of the major modes identified in the table require that the AP be informed. Transitions within the PSP mode (i.e. changes to the Listen_Interval) do not require any interaction with the AP. One situation in which mode transition may be appropriate is upon frame transmission, under the assumption that any transmitted frame is likely to be followed by a frame in the reverse direction. For example, a power-save station may transition to TAM mode following such a transmission, or alternatively such a station may operate in PSP mode while listening for all TIMs.

5.1.3 Access Point TIM Transmissions

Beacons containing a TIM (Traffic-Indication-Message) are sent at regular intervals. The TIM will identify the stations for which traffic is pending and buffered in the AP. In addition it contains an indication whether Broadcast/Multicast traffic is pending, and when it is going to be transmitted.

This information is coded in a *virtual bitmap*, as described in Section 6. To allow efficient coding, every station is assigned a Station ID code (SID) by the AP as part of the association process. The TIM will identify those stations for which frames are buffered by setting bits in the virtual bitmap which correspond to the appropriate SIDs.

5.1.4 TIM Types

Two different TIM types are distinguished. These are the TIM and the DTIM, which stands for Delivery TIM. Immediately after a DTIM, the AP will send out the stored Broadcast/Multicast frames, and all the frames stored (only those announced in the DTIM) for stations operating in the PSNP mode.

The DTIM interval is a manageable object, but would typically be selected to be between 50 and 200 msec. In between the DTIM's, TIM's can be scheduled at regular subintervals to increase the frequency of traffic announcements, typically between 20 and 50 msec.

The following figure illustrates the AP and station activity under the assumption that a DTIM is transmitted once every three TIMs.

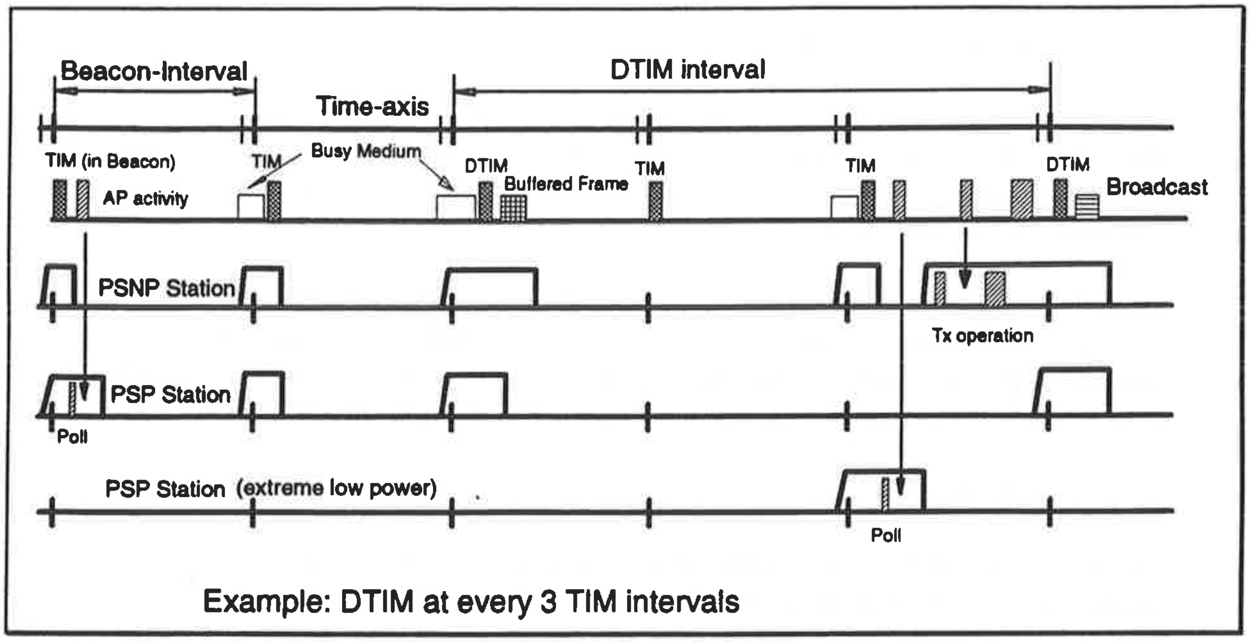

Figure 5-1: Infrastructure Power Management Operation

5.1.5 Access Point Operation

Access Points will maintain a Power Management status per station that indicate in which Power Management mode the station is currently operating. An AP will, depending on the Power Management mode of the station, buffer the frame destined to the station temporarily. All frames received for stations operating in the CAM or TAM mode are not buffered, and are directly forwarded. The station Power Management operational modes are indicated in the header of each frame transmitted to the AP. Stations can dynamically change modes, and indicate this in the frames transmitted to the AP, or they can use a separate NULL frame to change modes if necessary.

When an AP receives a frame to forward to the local BSS, then the Power Management mode status for the destination station has to be checked to determine whether the frame should be temporarily buffered, or whether it can be sent out immediately.

- Frames destined to PSNP and PSP stations need to be temporarily buffered in the AP.

- Frames destined to stations in the CAM or TAM mode can be directly transmitted.

- At every Beacon Interval, the AP will assemble the virtual bitmap containing the buffer status per destination for stations in the PSP and PSNP modes, and will send this out in the TIM field of the beacon. Also a Broadcast indication is included when Broadcast traffic is buffered.

- All Broadcast/Multicast frames will be buffered.

- After every DTIM, the AP will transmit the buffered broadcast/multicast frames, and the frames destined for stations in the PSNP mode.

- Buffered frames for stations in the PSP mode will be forwarded to the station after a Poll from that station.

- An AP will have an aging function to delete pending traffic when it was buffered for an excessive time period. The maximum age of a buffered item will depend on the Power Management mode (PSNP or PSP) of the station, or as further negotiated by the station at association time.

- The AP needs to maintain the Power Management status per associated station by monitoring the PM bits in the incoming frame headers.

- When an AP is informed that a station changes to the TAM mode, then the AP will send buffered frames (if any exist) to that station immediately.

DTIMs should be scheduled to allow adequate time to assure that packets are delivered within the timing constraints of LLC.

5.1.6 Receive Operation for Stations in PSNP Mode

Stations in PSNP mode operate as follows to receive a frame from the AP:

- Stations can choose to wake-up just before every DTIM and decode the virtual bitmap.

- When traffic is pending (either unicast or broadcast), then the station should stay awake after the DTIM to receive the buffered frames.

- A "More" bit in the frame header will indicate whether more data is pending, so that stations can stay awake until all buffered frames are received.

- When no traffic is pending stations can go into the Doze state again, scheduled to wake-up at the next DTIM.

5.1.7 Receive Operation for Stations in PSP Mode

Stations in PSP mode operate as follows to receive a frame from the AP:

- Stations can choose to wake-up immediately prior to any beacon. Since these stations need not hear every beacon (or the TIM information contained therein), they can be in the Doze state for a period longer than the beacon interval. The frequency with which a PSP station listens for beacons is a dynamic per-station parameter that depends on the power conservation and performance needs of the station, and could range, for example, from once every 20 ms (if beacons are that frequent) to many seconds.

- When traffic is pending (no broadcast), then the station should issue a Poll to retrieve the buffered frame.

- A "More" bit in the received frame will indicate whether more traffic for that station was buffered, so the station can Poll until no more frames are buffered for that station.

- To receive broadcast frames, the station should wake up at least before every DTIM.

5.1.8 PSNP and PSP Station Transmitter Operation

- When a station wants to transmit, it will wake-up from a possible Doze state, with sufficient time before the actual transmission is started, to ensure stable operation.

- A transmitting station will use the CSMA/CA access procedure, and transmit its frame(s).

5.1.9 Stations operating in the TAM Mode

- An AP will send down all the possibly buffered frames immediately after the station has indicated a mode change to the TAM mode to the AP.

- An AP will send incoming frames for stations in the TAM mode without any buffering delay.

- A bit in the frame header will indicate whether the station wants to change its mode to TAM, to change the AP buffering mode.

- A station shall switch to the TAM mode under control of a "Transmit_Holdover" parameter, which indicates how long a station stays awake after every transmission.

- When this value is non zero, then the station will change to the TAM mode, and remains in this mode under control of a "Transmit_Holdover" Timer.

- When this timer expires, then the station can, for instance, send a NULL frame indicating the mode change back to its original mode to the AP, and goes into the Doze state.

5.1.10 Stations Operating in the CAM Mode

Stations operating in this mode have their transceivers activated continuously, so they do not need to interpret the traffic announcement part of the beacons. Some stations want to synchronize to the beacon timer, for instance to allow support of Time Bounded Services.

5.2. Power Management in an Ad Hoc Network

In this section we show how the basic power management concepts for infrastructure networks described in Section 5.1 could be adapted to provide a similar service in the ad hoc case.

5.2.1 Basic Approach

The basic approach is similar to the infrastructure case in that the stations are synchronized, and frames which are to be transmitted to a power conserving station are first announced (in a quick transmission) by the transmitter. A power conserving station need only listen for these announcements (called Ad Hoc TIMs, or ***ATIMs***) to determine if its receiver must be left on.

The difference between the ad hoc and infrastructure case rests in the fact that there is no Access Point which is continuously powered and which can serve as a buffering and forwarding agent for the power conserving stations. Consequently each station must be responsible for generating its own announcements (i.e. the ATIMs). These essentially take the form of the header of the actual frame to be transmitted, including the source and destination addresses.

As was described in Section 4 above, the stations within an ad hoc network are synchronized through a distributed beaconing mechanism. One purpose of this synchronization is to allow stations to coordinate their awake times and the exchange of ATIMs.

The traffic announcement function is distributed among the transmitting stations and is a function of the Power Management mode of the destination station. When the destination station is in a Power Save mode, then a frame is announced by transmitting a ATIM frame by the transmitting station itself in a predetermined window in which all the stations operating in a power save mode are awake. In the case that a short frame is to be transmitted, then the frame itself is transmitted. The buffering function is thus essentially distributed over all the transmitters, in the stations themselves.

ATIM frames are only addressed to the destination station of the subsequent frame. In an Ad-Hoc Network only the CAM and PSNP mode are supported, while stations can under certain conditions temporarily switch to the TAM mode. PSP mode is not supported for Ad-Hoc networks.

5.2.2 Frame Reception with Power Conservation

A station which wishes to conserve power in ad hoc mode can operate as follows:

- The transceivers change to the "Awake" state at a predetermined "wake-up_window" time, triggered by the Beacon_Interval timer.

- Transceivers stay in the "Awake" for a predetermined interval, unless they receive an addressed ATIM frame (Station address or Broadcast).

- When stations receive a ATIM frame, then this means that the station should remain in the Awake state, because there will be a subsequent Data frame coming up.

This is illustrated in the following figure.

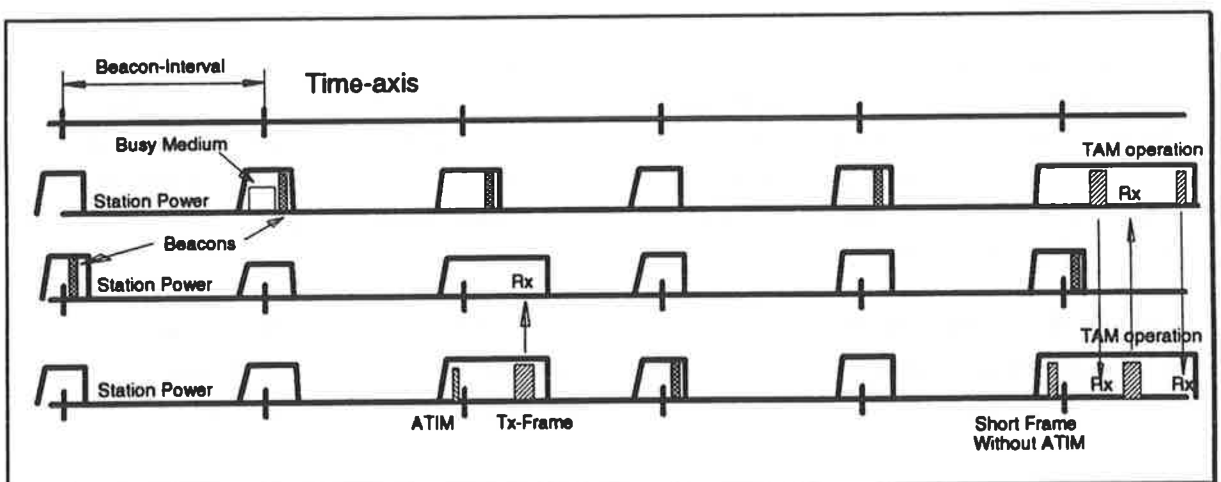

Figure 5-2: Ad Hoc Power Management - Basic Operation

5.2.3 Frame Transmission

Each station monitors the power-management status of the other stations. This is determined by examining the power-management bits within the frames generated by the other stations. This is actually only necessary for those stations with whom a given station needs to exchange frames.

To transmit a frame, the following steps transpire:

- When the destination station is operating in the CAM (or TAM) mode, then the station will transmit the frame without any buffering delay.

- Stations that have a frame to transmit to a destination station that is not operating in the CAM or TAM mode, will wait until their Beacon_Interval indicates a station awake status, which is some time (awake_Window) before the Beacon frame would be expected to arrive.

- Then this station will first put the destination station into the TAM mode by sending a short ATIM frame to the destination station, which is to be Acked by the destination station.

- The actual Data frame will normally not be sent until the awake period has elapsed, at which time only those stations that did receive (or transmit) a ATIM frame will stay in the Awake state.

- The timing of the start of the first ATIM frame should be randomized (backoff before transmission), to prevent multiple ATIMs from colliding.

- The ATIM is only the WMAC header of the Data frame, with the appropriate Type field set to indicate a ATIM.

Stations that have something to transmit to stations that are in the PSNP mode will incur an initial delay for the first packet, but it can then change to the TAM mode for a predetermined Transmit_Holdover_Time. This allows high speed data traffic directly between the two stations, once the initial connection has been established by sending the ATIM frame.

6. Contention Free Services

As discussed in Section 3, the DFWMAC does optionally support Contention Free services in certain environments. These services are provided by a PCF on top of a DCF, such that it assures coexistence with the Contention Services that are available in all stations.

This section will further specify the Contention Free services that are supported.
Two types of Contention Free Services are available:

- Time Bounded Service (TBS)
- Asynchronous Contention Free Service (ACFS)

The PCF will first handle the Time Bounded transfers in the TB-Burst, after which the remaining time available within the Contention Free limit can be occupied by the ACF-Burst. There is a big difference between the characteristics of both Contention Free services as will be explained below.

6.1 Time Bounded Services

This class of service is characterized by a relative long but bounded transfer delay, with low delay variance. The service is connection oriented, in a way that bandwidth must be reserved during a connection setup procedure. Different Quality of Service (QoS) levels may be needed.

For example, one of the potential applications is a voice connection that needs a very repetitive service, and allows some information loss. When the transfer delay reaches a given maximum then the information becomes obsolete, and the frame should then be dropped. Another application could be a Industrial Control Application, where the delay can be large, but must be bounded, while the information must be error free.

6.1.1 Time Bounded Frame Period

In the DFWMAC, the Time Bounded Service is to be reserved in terms of Time Bounded Frame Period (TBFP), and the maximum frame size per transfer. The TBFP must be an integer multiple of the SuperFrame Period (SFP). This is shown in the next figure.

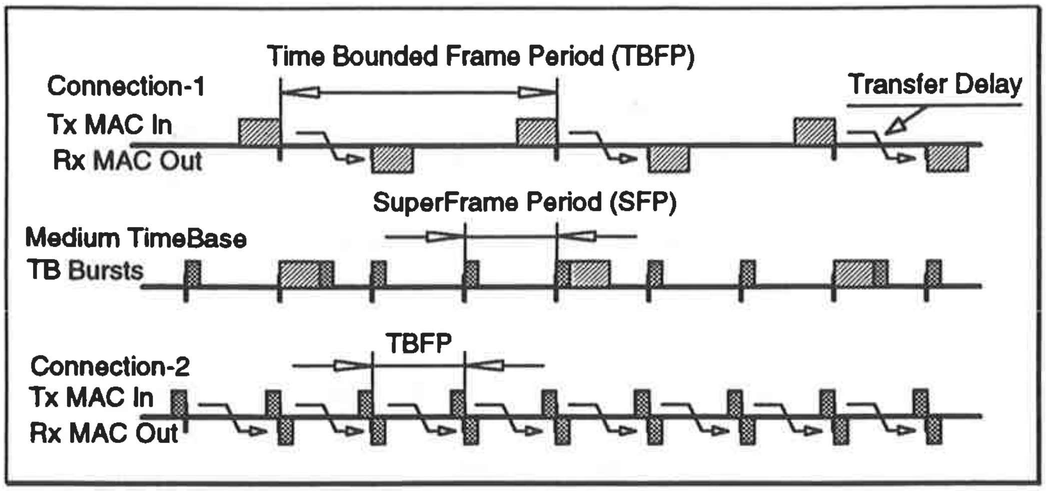

Figure 6-1: TBFP/SFP Relation

The Time Bounded Service is Connection oriented. The Asynchronous Contention service is used to setup a Time Bounded connection, which as explained will reserve a maximum frame size with a particular framing interval, the TBFP. When the request is granted, then the AP will give the station a Connection_ID. This ConnID is one of the fields that is used to address the station during the TB-Burst.

A separate frame structure is used for the Time Bounded frames. The main difference is that a 16 bit ConnID is used to address the station rather than the 48 bit Source and Destination addresses used for Asynchronous frames.

The format for the Type and Control Fields are the same for both Time Bounded and Asynchronous frames, but the rest of the frame has a different format. A bit in the Type field distinguishes between the two frame formats. Section 8 presents these frame formats.

6.1.2 QoS Aspects

Two different QoS levels could be distinguished. One can be a service level in which no recovery mechanism is used. Consequently, no Acknowledgment and doubling detection will be needed for this service.

Another QoS level would involve a limited recovery mechanism. The Ack provisions as provided by the PCF can be used to detect erroneous delivery. The actual recovery can be done in different ways. When the PCF detects an error in a TB-Down or TB-Up frame, then it can try to recover this within the TBFP period, so before the next frame for that connection is being generated. The PCF can do this for instance by extending the TB-

Burst with a retransmission of the TB-Down data for that connection (if there was a Down error), or with a TB-CTS to poll the station again to regenerate its TB-Up frame.

Another possibility could be to use the Asynchronous access scheme. The latter can be done by a station that never hears a Poll from the AP.

6.1.3 Time Bounded Service Implementation Characteristics

The described protocol has the following characteristics:

- Build on distributed access method under control of a Point Coordination function in the AP.
  - Uses CSMA/CA + (Ack) with highest priority.
  - Contains reservation mechanism to prevent Hidden stations from clobbering the contention free period.
  - Provides for limited Asynchronous recovery to increase robustness.
  - The connection establishment service uses the asynchronous data services.
- The protocol is dimensioned to support mixed Voice/Data.
  - Video support possible at higher PHY rates.
- Based on 20-25 msec SFP time (using a 2 Mbps PHY).
  - Shorter SuperFrame periods possible on faster PHY's.
- Support variable frame size on a per SFP basis without added control overhead. The station needs to negotiate the maximum size with the AP.
  - This allows to take full advantage of "Talk Spurt" characteristics of Voice.
  - It allows flexible congestion control for Voice.
  - No dummy data needed to fill unused reserved bandwidth.
- Unused reserved Time Bounded Bandwidth is fully available for Asynchronous traffic without any added control overhead.
- Can support different PHY speeds (1-20 Mbps)
- Includes provisions for Power Management and Re-association.
- Minimum burden for Asynchronous Data protocol implementation.

6.1.4 Time Bounded Service Capacity

Analyses have been done to evaluate the capacity of this system. The capacity of the system will depend on the PHY speed. The following results are based on a 2 Mbps raw speed of the PHY, using WAVELAN parameters where relevant, and a SFP of 20 msec. A 32 Kbps ADPCM voice source is assumed.

Because of the variable frame size capability, it is easy to deal with a Talkspurt Voice System, that will only generate data as function of a voice activity indicator (use factor of .4 is assumed).

The results are calculated assuming a MAC frame overhead of 15 Bytes and the PHY overhead using the WAVELAN parameters.

The results are as follows:

Voice only:	28 Full Duplex channels (using Talkspurt characteristics)
Voice/Data 576 Bytes:	23 Full Duplex channels (using Talkspurt characteristics)
Voice/Data 1500 Bytes:	18 Full Duplex channels (using Talkspurt characteristics)
Voice/Data 1500 Bytes:	9 Full Duplex channels (assuming 100% load)

Asynchronous Data throughput:	> 75 KByte/sec (assuming max TBS load)
Asynchronous Data throughput:	>130 KByte/sec (assuming 9 FDX Talkspurt channels)
Asynchronous Data throughput:	>200 KByte/sec (no TBS connection active)

In the above example, (given the 1500 Byte max Asynchronous frame size) the maximum number of connections per cell can best be limited to 9, which still allows for a minimum of approx. 75 KByte of Asynchronous data throughput capacity up to more than 130 KByte throughput capacity when Talkspurt Voice characteristics are taken into account.

The SFP interval and so the absolute MAC-to-MAC delay is dependent on the PHY speed and the maximum frame size of the Asynchronous Data.

The efficiency will increase for longer SFP intervals.

6.2 Asynchronous Contention Free Service (ACFS).

Asynchronous traffic is characterized by its bursty nature. The ACFS is designed to efficiently use the available bandwidth while dealing with its bursty nature. The MAC can not predict in advance how much bandwidth is needed at any time to serve the application. Therefore it can not reserve any bandwidth ahead of time.

Instead the PCF can dynamically allocate time to a station that requested Contention Free service. In the AP all down traffic that is queued up for transmission at the end of the TB-Burst, or even during the ACF-Burst can use the PCF to allow Contention Free medium access for its buffered traffic.

If a station wants to use the ACFS, then it needs to request for this service using the Contention period. This does not reserve any bandwidth, it just signals its capability to the AP, so that the AP can when the need is there, put the station in its Polling list.

Stations (that did an ACFS service request before) can do different things to get service in the contention free period.

- They can do a Poll request in the Contention period.
- Or, they can send their frames in the Contention period. This can be detected by the AP, which can then notify the PCF to put the station in the Polling list.

The PCF should monitor the station activity in the Contention Free period. When the station did for instance not respond to a Poll for a number of consecutive times, then the PCF could delete the station from its Polling list to prevent that more bandwidth is wasted in idle polling.

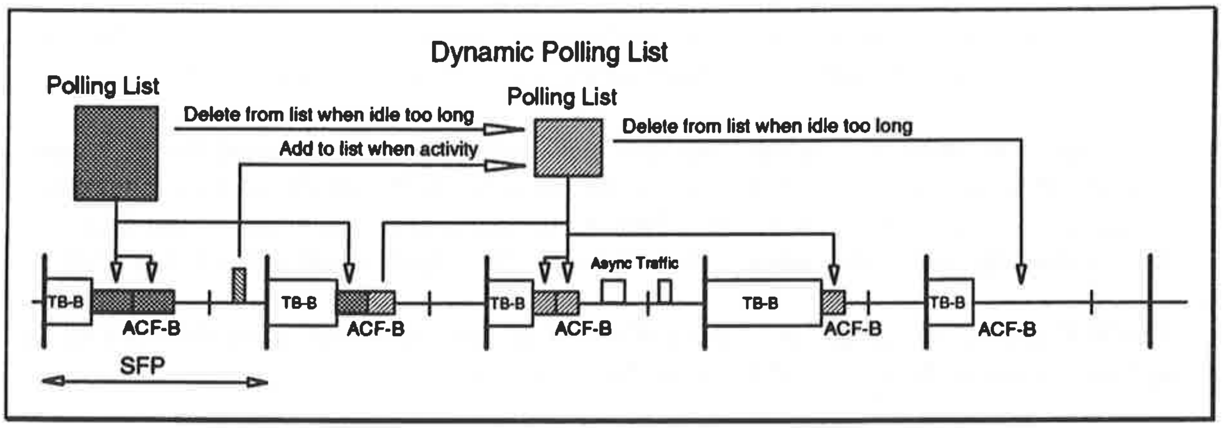

Figure 6-2: Low Load Scenario

The frame format for the ACFS is the same as the normal Asynchronous Data frame. The Poll bit is part of the fixed header of the frame. The PCF in the AP will transmit a buffered Asynchronous frame to the station with its Poll bit activated. Or when no data is buffered, but the station is in the active polling list of the PCF, it will send an empty data frame addressed to the station with the Poll bit on.

Note that the RTS/CTS option should not be used in the Contention free period, because no contention resolution is needed when the PCF is used.

Stations can still use the Contention period to transfer their data when needed.

6.2.1 ACF-Burst

The PCF will after the TB-Burst generate an ACF-Burst when there are entries in the Active Polling List. It will work through the polling list, generating a Poll to the stations, until the Contention Free limit is reached, or until the full polling list is cycled n times within the same SuperFrame period.

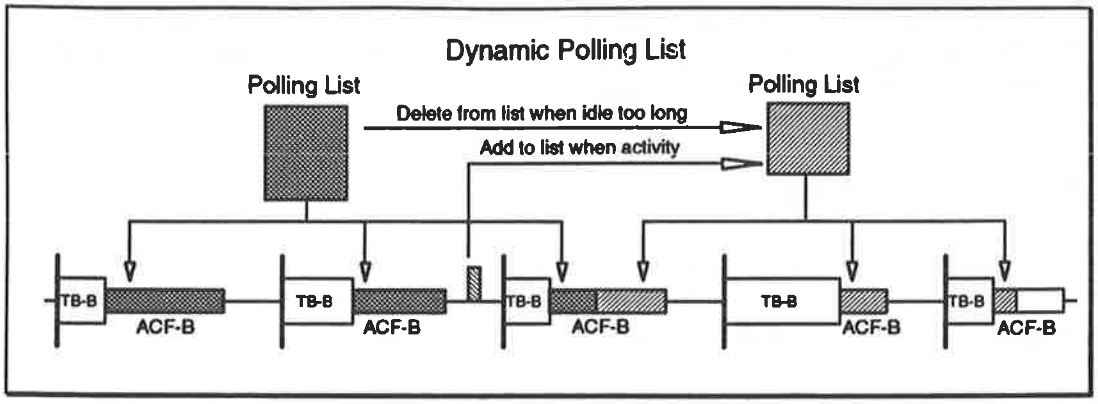

Figure 6-3: High Load Scenario

If the Poll list is longer then can be serviced within the maximum Contention Free period of the SuperFrame, then the PCF should continue in the next SFP. This means that within a network with many active stations, then it can take several SuperFrame periods to service every station in the list.

6.2.2 ACFS Characteristics

If stations want to use only the ACF Service then they will get contention free access, but at the expense of a higher access delay. This can be advantageous during heavy load periods, but will increase the access delay also during the lower load situations.

If stations use both, then the PCF can automatically take control during high load situations by servicing the stations also in the Contention Free period. This can increase system capacity under certain circumstances.

7. MPDU and Frame Formats

The basic unit of information which is exchanged between MAC entities is a *frame*. A complete MAC Protocol Data Unit (MPDU) may consist of a sequence of related frames exchanged between two MAC entities (for example, RTS/CTS/DATA exchanges or DATA/ACK exchanges). In such a case, the relationship among the frames is indicated via an MPDU ID field within the frame headers.

7.1 Basic Frame Format

Each frame consists of the following basic components:

- A fixed-length *PHY-Adaptation Header* (PHY-dependent)
- A fixed-length *Fixed Header*, which includes a Type field, Control field, and MPDU-ID.
- A variable length *Frame Body*, including

  Address fields, which may include one or more fields with Network ID, Destination Address, and/or Source Address, depending on the frame's type.

  Variable length components, including a (possibly empty) set of optional *elements* and a (possibly empty) set of *type-dependent fields.*

- An 8-bit or 32-bit CRC, depending upon the frame type.

Frames may optionally include a variable length "Elements" subfield which contains one or more "elements". Elements are defined to have a common general format consisting of a one-octet Code field, a 1 bit More indicator (identifying whether additional elements are present), a 7-bit Link field, and a variable-length element-specific field. Each element is assigned a unique code. The Link field is the number of additional octets in the element (which may be zero).

Certain frame types may require that specific elements be present.

The figure below depicts this general format:

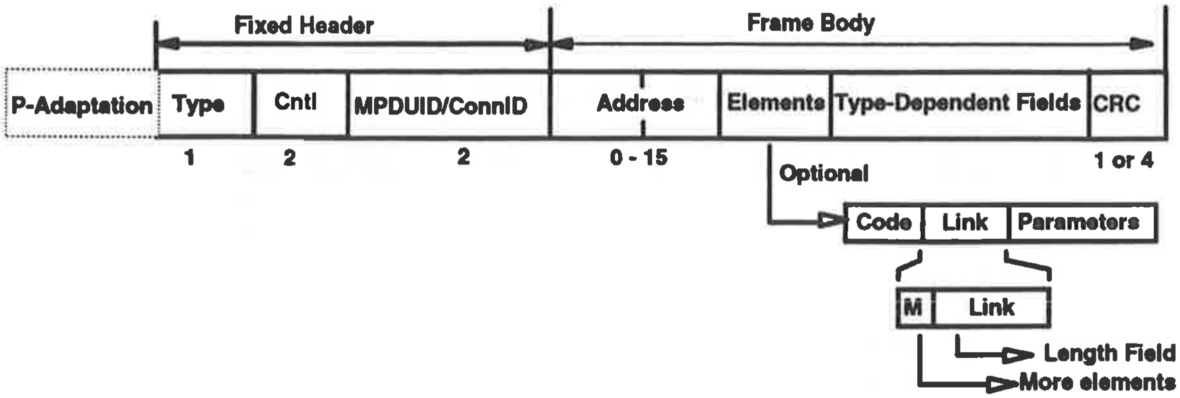

Type Field Format

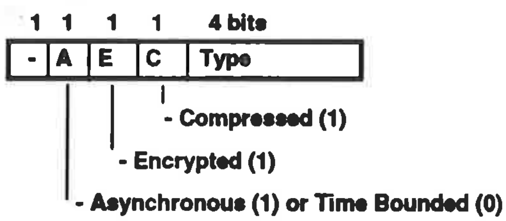

NID Format (part of Address within certain frame types)

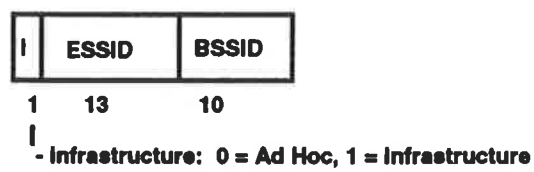

Control Field Format

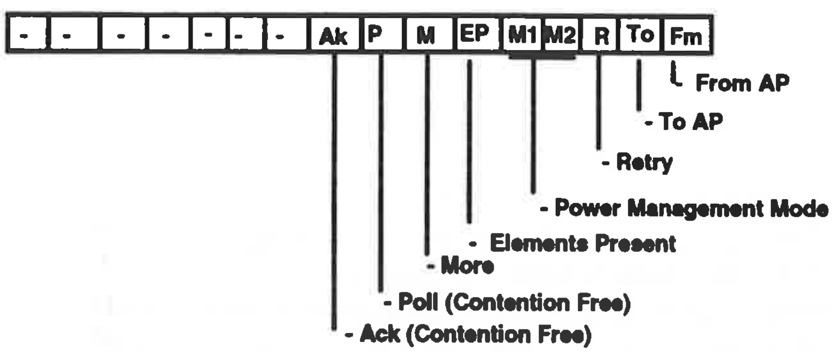

The fields depicted in this figure are described below:

P-Adaptation

A PHY-dependent, fixed length field.

Type

1 octet, including a 4 bit type field and 3 control bits. The control bits indicate asynchronous or timebounded service, whether or not the frame is encrypted, and whether or not the frame is compressed. The type subfield takes the following values:

Asynchronous
- 0 Reserved
- 1 RTS (Request to Send)
- 2 CTS (Clear to Send)
- 3 Data
- 4 Ack
- 5 Poll
- 6 Beacon
- 7 ATIM (Ad-hoc Traffic Indication Map)
- 8 Request
- 9 Response

Time Bounded
- 0 Reserved
- 1 TB-Up
- 2 TB-Down
- 3 TB-CTS
- 4 Ack

Control

2 octets, the bits of which have the following meaning:

Ack: for contention free service acknowledgment
Poll: for contention free service poll
More: indicates that additional data frames are buffered
Elements Present: indicates that optional elements are in the frame
Power Management Mode: 2 bits, as follows:
- 00 CAM
- 01 PSP
- 10 PSNP
- 11 TAM

Retry: 1 = this is a retransmitted frame, 0 = new frame
To: 0 = to access point, 1 = to station
From: 0 = from access point, 1 = from station

MPDUID/ConnID

This 16 bit field is the MPDU-ID for asynchronous frames and the Connection ID for time-bounded frames. The MPDU-ID value is the hash of the 2 octet NID field, 6 octet source address and a 1 octet sequence number. Time-Bounded ConnectionIDs are assigned when the station associates in a time-bounded service mode.

Address

Type-dependent length and format, may contain the 16 bit NID and/or 48 bit IEEE addresses (Destination, Source, neither or both, depending upon frame type). The NID is a 3-octet NetworkID field, consisting of a 1-bit "Infrastructure" indicator (0 denotes Ad Hoc, 1 denotes Infrastructure), a 13 bit ESS ID field, and an 10 bit BSS ID field.

Elements

Variable length field containing zero or more elements.

Element format includes 1 octet Code field, 1 bit More indicator, and 7 bit Link field.

Link field indicates the number of remaining octets in the element.

Type-Dependent

Additional variable-length type-dependent fields may be present.

CRC

1 octet (CRC8) or 4 octets (CRC32).

The following figure displays the type-dependent fields within the various frame types.

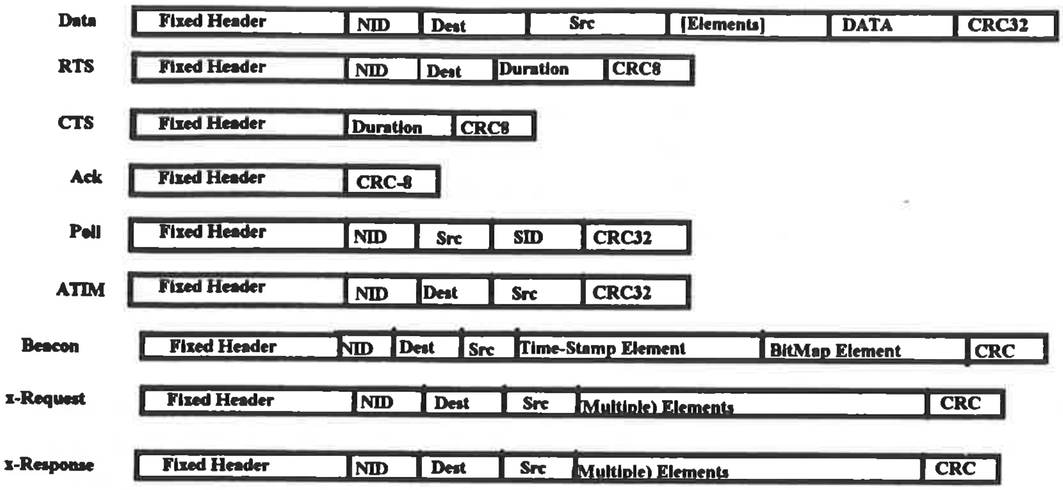

The following Request Types are defined:

- Associate
- Probe
- Null
- Beacon

7.2 Element Definitions

The general format of all elements is depicted in the following figure:

General Element Format

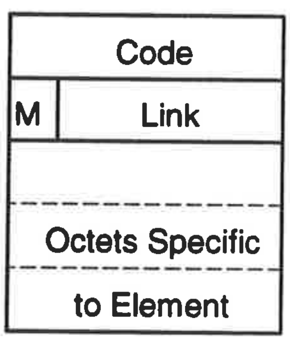

More Bit M = 1: More Elements Follow
= 0: Last Element

Link (7 bits): number of subsequent octets

Defined elements include the following:

Timestamp

This element is 4 octets long. The value represents the current clock value of frame's source in microseconds. The high order bit, if set, indicates that the station is synchronized within its BSS.

Beacon Interval

1 octet, representing number of milliseconds between Beacon generations.

Traffic Indication Map (TIM)

Variable number of octets. The TIM element contains a variable number of *block groups*, with each block group consisting of a *block identifier* followed by 0 to 7 one-octet *blocks*. Each bit within a block indicates whether a frame is currently buffered for a station with a particular Station ID. There is a one-to-one mapping between the bits in a *virtual bitmap* and the station IDs which is maintained within the access point; the actual transmitted TIM is a compressed representation of the virtual bitmap.

Block (8 bits)

Each bit corresponds to a specific station within the block. If this block represents the Nth block within the virtual bitmap, then Bit M within the block corresponds to the station with Station ID equal to 8*(N-1) + M.

Bit = 1: There is a frame pending for this station
Bit = 0: There is no frame pending for this station.

BI: Block Identifier (1 octet)

M	B1	B2	B3	B4	B5	36	B7

Bit 0: More   0 = This is the last block group
               1 = Another block group follows

Bit N (N = 1..7)   0 = Nth block in this group is absent
                    1 = Nth block in this group is present

Block Group: Consists of a Block Identifier followed by 0 through 7 Blocks.

BI	Blocks
Block Identifier	0 - 7 blocks

DTIM Count

One octet. Indicates how many TIMs (including the TIM in the current frame, if any) will appear before the next DTIM. A DTIM Count of 0 indicates that the current TIM is in fact a DTIM.

DTIM Period

One octet. Indicates the number of TIM intervals between successive DTIMs. If all TIMs are DTIMs, the DTIM Period element will have value 1.

Broadcast Indicator

Zero octets. If present, this element indicates that a broadcast or multicast frame will be transmitted by the Access Point following the next DTIM (or after the current frame if this frame includes a DTIM).

Station ID

Two octets, representing the 16-bit Station ID of a station.

Associate

Zero octets. If present in a REQUEST frame generated by a station, indicates a desire on the part of the station to associate with a given access point. If present in a RESPONSE frame generated by an access point, indicates that this is a response to a prior association request.

Weight

16 bit value indicating weight of station. Used within Beacons and Probe Responses.

Source NID

Represents source NID of transmitting station. Used in Probe Responses.

7.3 Elements Within the Frame Types

The formats of DATA and ACK frames have already been discussed. The formats of the other frame types consist of the basic format (defined above) with specific values within the Type field. The various types may also include elements as described below.

POLL

The *StationID* element is mandatory.

ATIM

There are no mandatory elements for an ATIM.

BEACON

The *Timestamp* elements is mandatory. The *TIM, DTIM Period, Weight*, and *Broadcast* element may also be present.

REQUEST

If an *Associate* element is present, then the *Previous AP Address* element must also be present.

RESPONSE

If an Associate element is present, then the *Timestamp*, *Station ID*, *DTIM Period*, and *Beacon Interval* elements must also be present. A Probe Response includes *Weight*, *Timestamp*, and *Source NID*.