IEEE 802.11
Wireless Access Method and Physical Layer Specifications

Title: **The CODIAC Protocol**

Centralized Or Distributed Integrated Access Control (CODIAC), A Wireless MAC Protocol

Author: Carolyn Heide
Spectrix Corporation
906 University Place
Evanston, Illinois 60201
Voice: (708) 491-4534
FAX: (708) 467-1094
EMail: 71041.3262@compuserve.com

Abstract:

This paper proposes a media access control (MAC) protocol to address the diverse requirements of the IEEE 802.11 standard effort.

The goals of the proposed protocol are: (1) to take advantage of the contention avoidance, power efficiency and time-bounded service support characteristics of a deterministic MAC protocol; (2) to operate with efficiency and fairness in the absence of infrastructure; and (3) to provide maximum flexibility, allowing the protocol to be tailored to varying implementations without loosing compatibility across implementations.

Related Documents:

IEEE 802.11-93/40 The WHAT MAC Protocol
IEEE 802.11-93/59 Evaluation of CODIAC Protocol
IEEE 802.11-93/65 Performance of CODIAC Protocol

1. Introduction

This protocol stems from a merging of two very different protocols, each of which is ideally suited for a particular wireless LAN scenario, but not for all scenarios.

The first protocol of influence is the Spectrix Reservation/Polling Protocol[1], which is currently implemented in high density, fixed population networks - specifically, 1500 mobile stations per Basic Service Area (BSA). This implementation of the protocol is described briefly in Appendix A. It is a deterministic collision avoidance, master/slave protocol. No station may access the medium without permission, implicit or explicit, of the master and so the protocol is collision free despite the large number of stations and the possibility of hidden stations. However, the strict deterministic nature of this protocol makes the issues of adhoc networks and BSA overlap difficult to handle.

The second protocol is the WHAT protocol proposed by XIRCOM in several IEEE 802.11 submissions[2], described in detail in document IEEE P802.11-93/40. This protocol excellently handles wireless issues, such as hidden stations, overlapping service areas, and ad-hoc networks. It is simple and elegant, and has been implemented and simulated with success for small station populations with office LAN traffic patterns. However 802.11 must address scenarios of more varying traffic patterns and much greater station density, and the ability of an Enhanced Listen-Before-Talk protocol to do so without major performance degradation is questionable.

The protocol proposed here endeavors to combine the best features of the these two protocols, resulting in a protocol with the flexibility to support the diverse requirements of the 802.11 standard.

[1]Patent Pending, US Serial No. 07/643,875

[2]IEEE P802.11/91-92 A Hybrid Wireless MAC Protocol; IEEE P802.11/92-14 An Update to the Hybrid Wireless MAC Protocol; IEEE P802.11/92-49 A Review of Some Properties of the Hybrid Protocol; and, IEEE P802.11/93-40 The Wireless Hybrid Asynchronous Time-bounded MAC protocol.

2. The Concept

Take the WHAT protocol without the T - the asynchronous support only. All data transfer is via a four step frame sequence which is issued when the Net Allocation Vector (NAV) shows a gap in the bandwidth allocation:

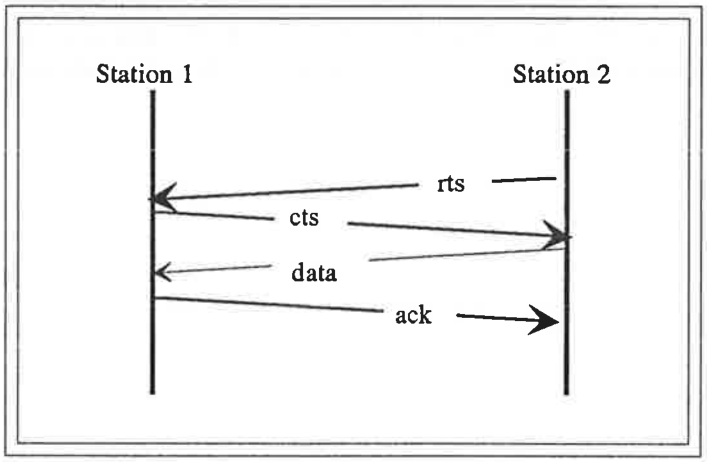

Figure 2.1

Then take a Reservation/Polling Protocol (RPP) where all data transfer is via the four step frame sequence that is initiated by a synchronization frame from a controller:

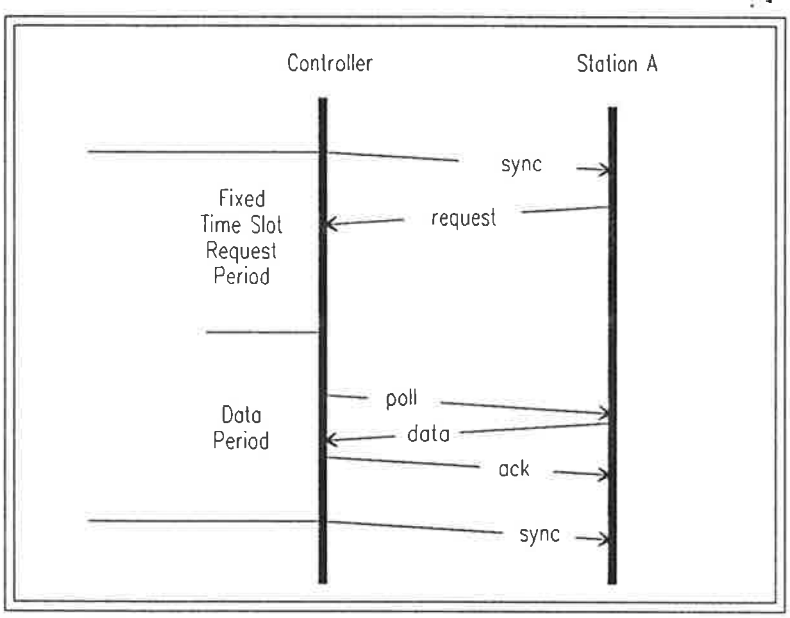

Figure 2.2

Combine the processes:

- Change the REQ frame to an RTS frame - a request for bandwidth. The request specifies the destination of data. If there is a controller present, send the request to it in response to the

SYNC frame and let it allocate the bandwidth. If not, broadcast the request according to the NAV;

- Change the POLL to a CTS - the controller is not asking for data from the station, but clearing the bandwidth requested for the station to send it's data;
- The DATA and ACK follow the CTS just as in both the WHAT and the RPP.

To a station the only difference between the two protocols is when the RTS is issued - either according to the NAV, for following a SYNC frame from a controller. Stations could easily have two modes of operation according whether or not they are in the BSA of a controller, and these modes would differ very little.

3. Theory of Operation

The CODIAC protocol has two modes of operation:

(1) Centralized: In the presence of a controller station, which is issuing a periodic, but not necessarily regular, synchronization frame. In centralized mode access to the medium is highly deterministic, and completely controlled by the controller station. This mode of operation supports both 'asynchronous' and 'time-bounded' service;

(2) Distributed: In the absence of a controller station, when a station cannot 'hear' sync frames it operates in distributed mode, enhanced Listen-Before-Talk (LBT). This mode of operation does not support time-bounded service.

It is not required that the controller station be an Access Point (AP), or that an AP must be a controller station. Nor is either mode associated with an adhoc or an infrastructure network - an adhoc network could be created by bringing a portable controller station into a room, or an infrastructure network with no requirement for time-bounded services could function in distributed mode.

The startup procedure for a station, when it is powered up or comes 'awake' is:

(1) Listen for x amount of time for a sync frame (where x is an known maximum time);

(2) If sync heard, operate in centralized mode;

(3) If no sync heard, operate in distributed mode.

In either mode of operation the transfer of data is accomplished by a four step transaction - the exchange of RTS, CTS, DATA and ACK frames. The originating station sends a Request To Send (RTS) frame. In response to the RTS, a Clear To Send (CTS) frame is sent. On receipt of the CTS frame the originating station sends the DATA frame, and the destination responds with an ACK frame. The timing of the transmission of these frames differs between centralized and distributed mode, but not their format or meaning.

4. Frame Format

All frames have the following format:

Minimum Frame Length (12 +n) octets								
		FCS Coverage						
Preamble	SD	DID	Type	Control	INFO	FCS	ED	⇐ field name
8n	8	16	8	8	(optional) 8m	32	8	⇐ field length (bits)

Preamble	= Preamble (n to be determined)
SD	= Start Delimiter
DID	= Destination Identifier
Type	= Frame Type
Control	= Control Flags: AP, sequence, out-of-sequence, retry, hierarchical
INFO	= Information (0 < = m < = to be determined)
FCS	= Frame Check Sequence, CRC-32
ED	= End Delimiter

4.1. Destination Identifier

The purpose of the destination identifier (DID) field in the frame is to allow the MAC to determine as quickly as possible whether this frame needs to be received by this station. That does not mean that it contains the same value for a specific station at all times.

If a station is registered with a controller, the controller has assigned an ID to that station. The controller will use this ID as DID in all frames it sends to that station, and it will ensure that all stations registered have exclusive IDs.

When stations are not registered with a controller (i.e. operating in distributed mode), they choose any value for their ID, and may vary it with each transaction (this is equivalent to the WHAT protocol's MSDUID, which is incremented with each MSDU sent).

Destination identifier values:

(1) FFFFh = broadcast;
(2) 8000 - FFFEh = controller stations;
(3) 0 - 7FFFh = non-controller station.

4.2. Control Flags

AP - indicates frame originated from an Access Point.
Sequence - alternation bit, one bit sequence number (see section 8).
Out-of-sequence - indicates invalidity of sequence bit (see section 8).
Retry - indicate re-transmission (see section 8).
Hierarchical - specifies frame destination must be Access Point only.

5. Distributed Mode Operation

Distributed mode of operation is the asynchronous operation of the WHAT protocol. The CODIAC protocol uses slightly different frame formats, but the operational rules are the same.

When operating in distributed mode all stations must have their receivers on at all times if: (1) there is any possibility of them receiving data; or (2) they wish to send anything.

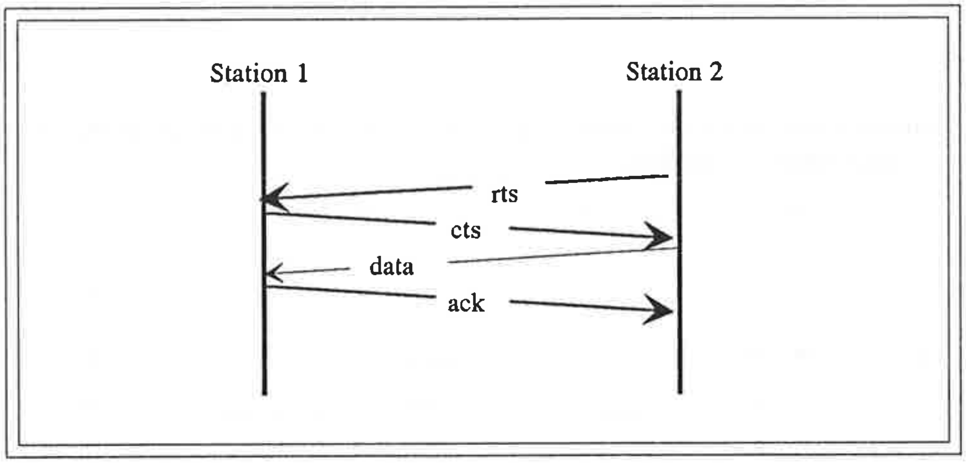

Figure 5.1 - Distributed Mode: Data from Station 2 to Station 1

6. Centralized Mode Operation

6.1. Introduction

The characteristics of the centralized mode of operation allow a lot of flexibility for individual implementations. The following sections cover the most basic possible method of operation. In later sections some possible enhancements (i.e. complexities!), and the reasoning behind them, are introduced.

6.2. The Superframe

Stations operate in centralized mode when they know they are in the presence of a controller station because they receive synchronization frames from it.

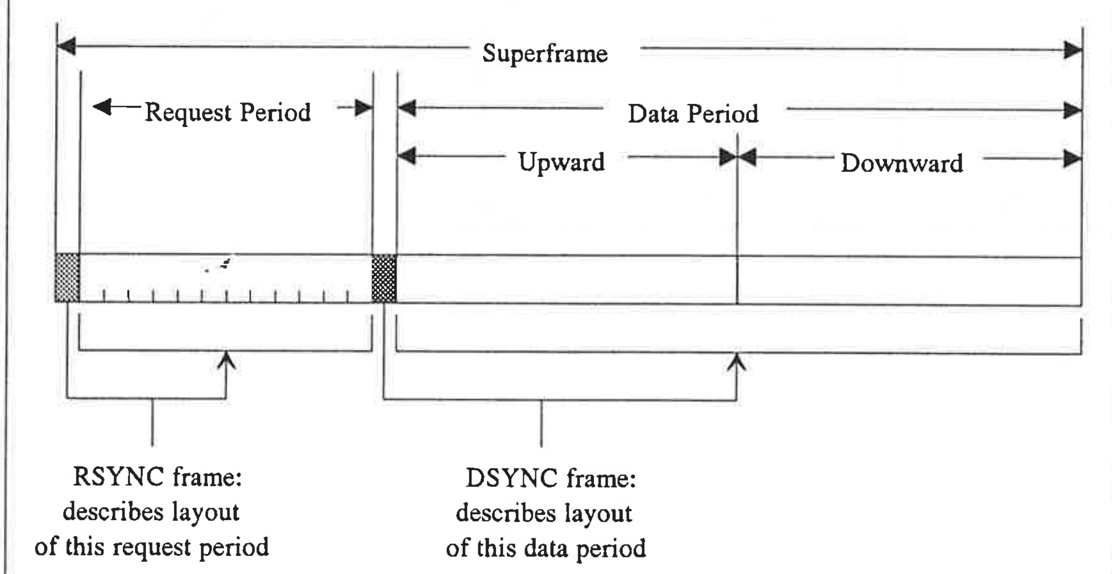

Figure 7.1 - Superframe Layout

Centralized mode operation takes place in quantums of time called superframes. Each superframe has two periods - the request period and the data period. The length of each of these periods is dynamic and controlled by the controller station. Each period is initiated by transmission of a synchronization frame by the controller - an RSYNC frame marks the start of the request period, and a DSYNC the start of the data period.

The request period is divided into fixed length time slots which are used by stations to register with the controller, and to request bandwidth to transmit after they have become registered. During this period stations will transmit only, nothing will ever be sent to them, so stations may turn off their receivers for this period. The purpose the RSYNC frame at the start of the request period is to facilitate synchronization of the time slots, and to specify the length of time for which stations may turn off their receivers.

The data period is itself divided into two parts: upward (data sent from stations to the controller) and downward (data sent to stations, either from the controller, or from other stations). The purpose of dividing the data period is to allow stations which have no data to send to the controller - those which have not issued requests in the request period - to turn off their receivers during the upward data period. The purpose of the DSYNC frame sent by the controller at the start of the data period is to specify the length of the upward data period.

During the upward data period stations which requested to send data to the controller will be allocated bandwidth by the controller to do so. Stations which did not request to send to the controller during the request period may keep their receivers off, they will not be sent anything during this period.

During the downward data period data is sent to stations - either from the controller, or from other stations which requested direct station-to-station transmission during the request period. During this period all stations must keep their receivers on (unless they have some way of guaranteeing that no one is going to send them data).

6.3. Request Period

The request period is divided into fixed time slots. The length (in bit times) of each slot is fixed, while the number of slots for each sync period is specified by the controller and placed into the RSYNC frame. There are two types of time slots, registration slots and owned slots. Once a station has registered with the controller, the controller assigns it a time slot, which it then owns.

A station looses ownership of a time slot when:

(1) The station cancels its registration with the controller;

(2) The lifetime of the ownership, which is a time length known to both the controller and station, expires without the slot having been used; or

(3) The controller cancels the registration of the station.

To register with a controller, a station generates a random number and uses this to determine in which registration slot to send its registration request. Should it collide with another registering station, it's request will not be answered by the controller, and it will repeat the process next superframe.

While the registration request is a special frame type, all other registration related information is passed in MAC management data frames (MDATA). This is so that registration acceptance, rejection and cancellation information can take advantage of the acknowledgment that goes along with data transactions. Information about time slot ownership must be protected against loss so no confusion over ownership can arise.

After a station has issued a registration request in a registration time slot, the controller responds with a Registration Accept contained in an MDATA frame. The MDATA frame is sent to the ID specified by the requesting station in the registration request. The Accept specifies the number of the time slot which is now owned by the station. From this point onwards, the ID of the registered station is it's time slot number.

The total number of registration slots available in each superframe is determined by the controller and specified in the RSYNC frame.

The owned time slots are the first time slots following the RSYNC frame, and the registration slots follow these. For example, if the RSYNC specifies 1024 total slots with 100 registration slots, slots 1 through 924 must be owned, and slots 925 through 1024 are available for registration requests. If a new station registered the distribution of slots in the next request period is at the discretion of the controller - there can to 1025 total slots with 100 registration slots, or 1024 total slots with 99 registration slots, or any other combination as long as there are enough slots for the stations which own them.

When an owning station gives up, or looses, its time slot, this slot is now available to the controller to assign to the next station which registers. Thus the 'holes' in the owned time slots will be filled up, and no need is foreseen to have to re-pack the slots. Should some extreme condition cause a very poor distribution of owned time slots, the controller can de-register stations, and when they register again assign them to the empty time slots.

It should be noted that registration with a controller is not equivalent to association with an Access Point. The purpose of registration is to facilitate coordinated use of the medium by all stations within range a controller. For instance, two stations may be registered with a controller/AP so that they can converse with each other without disrupting communication in the BSA, but they may not be associated with the AP and the infrastructure at all. They are forming a small adhoc network in the presence of the AP's infrastructure network, but both networks are using a common point coordination function under control of the controller.

Special Cases:

At the discretion of the implementer, the owned slots and the registration slots could be the same. If the RSYNC frame specifies this, then requesting stations are sharing their request slots with registering stations. Any time a station has data to send it has the possibility of having its request collide with a registering station - this may be a risk some implementations find acceptable. Since the registering station chooses its slot based on a random number, collisions it causes should not seriously delay transfer of data from any particular station, and the odds of registering in a slot that is not about to be used by its owner can be very high, depending on the application.

6.4. Registration Transaction

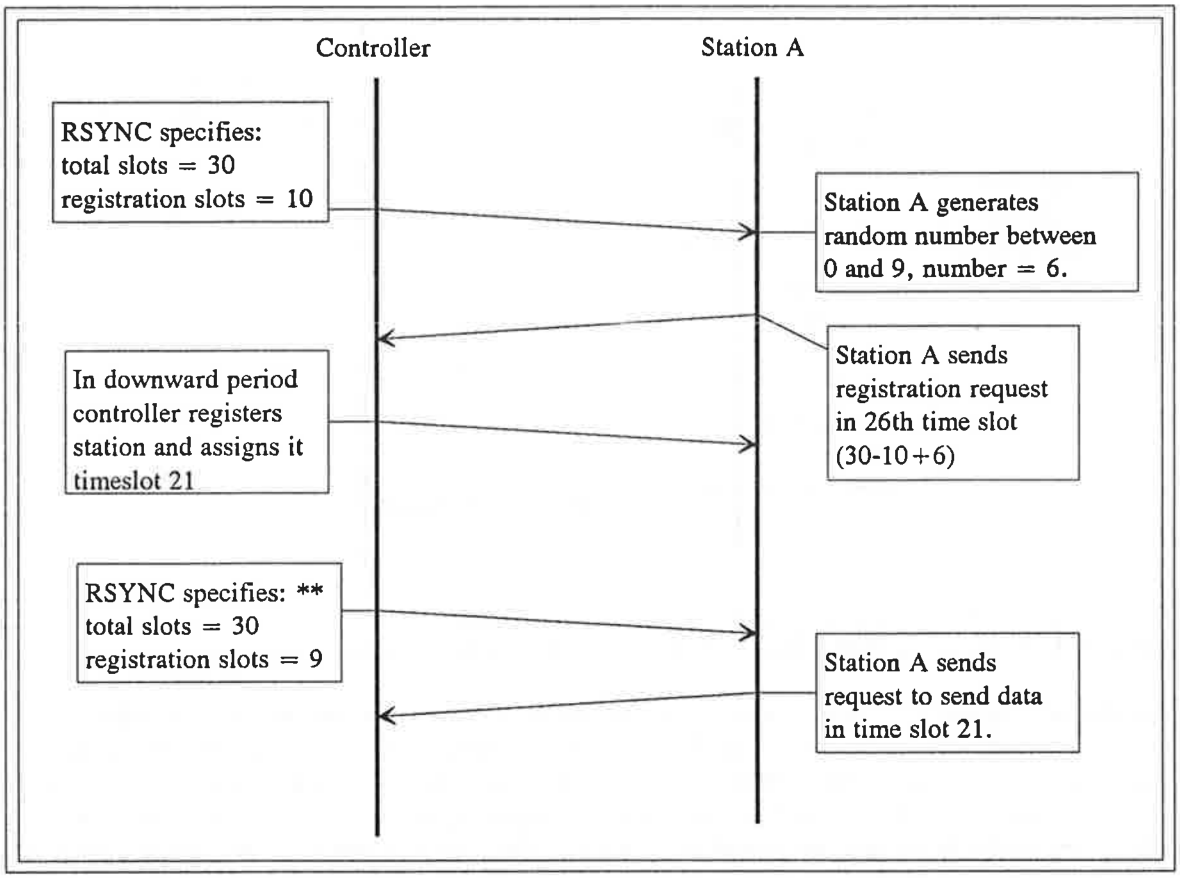

Figure 7.2 - Station A registration. ** the next RSYNC after station A registers could contain total slots 31, registration slots 10, if the controller wants to keep the number of registration slots fixed.

6.5. Upward Data Period

At the end of the request period the controller has a list of all stations which want to transmit, and the destination and length of each transmission. The controller then initiates the data period by broadcasting a DSYNC frame. The DSYNC frame begins the upward data period, the length of which is specified in the DSYNC frame. The controller allocates the bandwidth within the upward data period by sending a CTS frame to, and receiving a DATA frame from, each station from which it wishes to receive data.

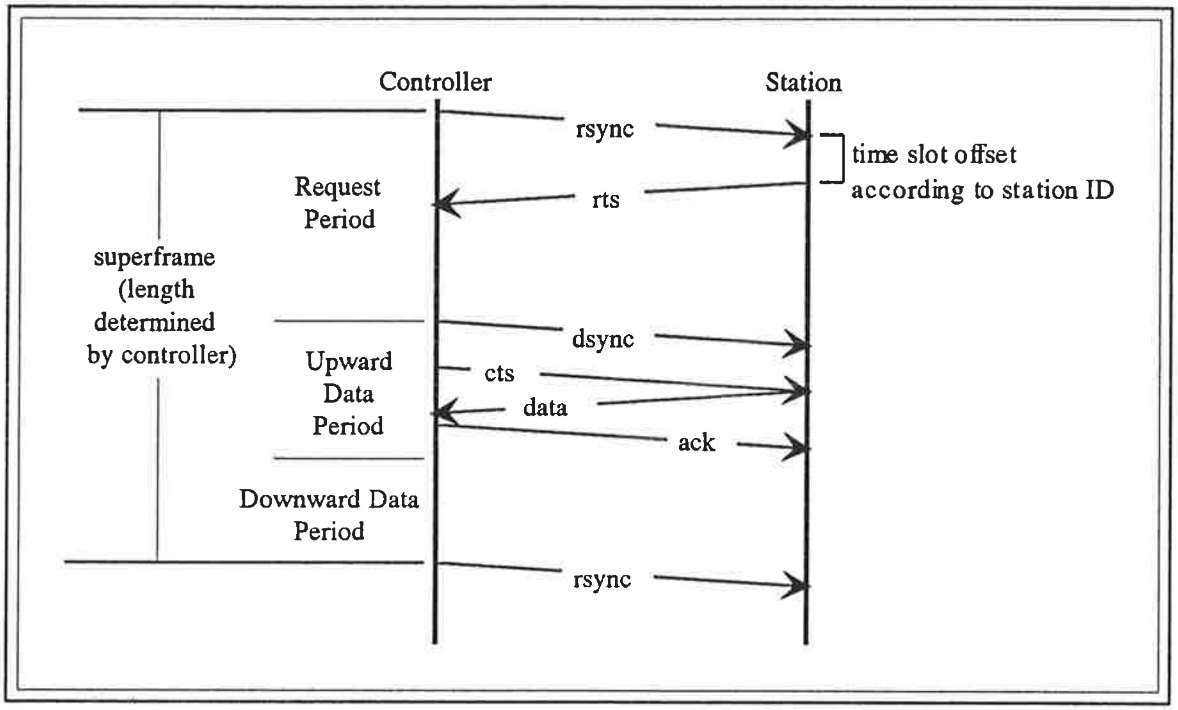

Figure 7.3 - Centralized Mode: Data from Registered Station to Controller

The station's original RTS, although sent to the controller, may not have specified the controller as the destination of the data to be sent. This destination was specified in the RTS by the unique 48-bit address of the destination station. If the controller knows that the destination station is a registered wireless station, the CTS frame will not be sent to the requesting station in the upward data period. It will be sent in the downward data period because that is the only time when the destination station is guaranteed to have its receiver enabled.

There could exist a case when a station's RTS specified a destination which is a registered wireless station, but although both stations are in range of the controller, they are not in range of each other. For this reason, the sending station can specify that the controller should send the data indirectly, by using an RTSI frame rather than an RTS frame. When the controller sees an RTSI frame, it specifies in the CTS that the data should be sent to the controller even if the data destination station is registered. Once the controller receives the data it queues it to be sent as downward data to the destination station, much as if it had originated outside the wireless network.

6.6. Downward Data Period

During the downward data period data is sent to non-controller stations, from the controller and from other stations. This activity is restricted to this period to aid in power consumption, because it requires that all stations keep their receivers on.

In this period the controller sends to stations data which it has for them from other stations in the BSA which requested indirect station-to-station data transfer. If the controller is an AP, it also sends to stations data it has received for them from the distribution system.

There are two possible methods by which the controller can send data to a station during this period, and implementations of the protocol may choose to use either or both:

1. The controller may just send a DATA frame to the station. In this case the controller risks wasting bandwidth if the station is not present. It also risks having stations which overlap with the BSA and are functioning in distributed mode cause loss of the data due to collision;
2. The controller can send an RTS first, get a CTS back from the station and then send the DATA frame. This will save bandwidth in the two cases described above, but cost bandwidth in most cases.

An effective implementation may be to send the DATA frame alone the first time, and if a retry is required then to use the full handshake method. Or to revert to the full handshake based on an overall failure rate of some kind.

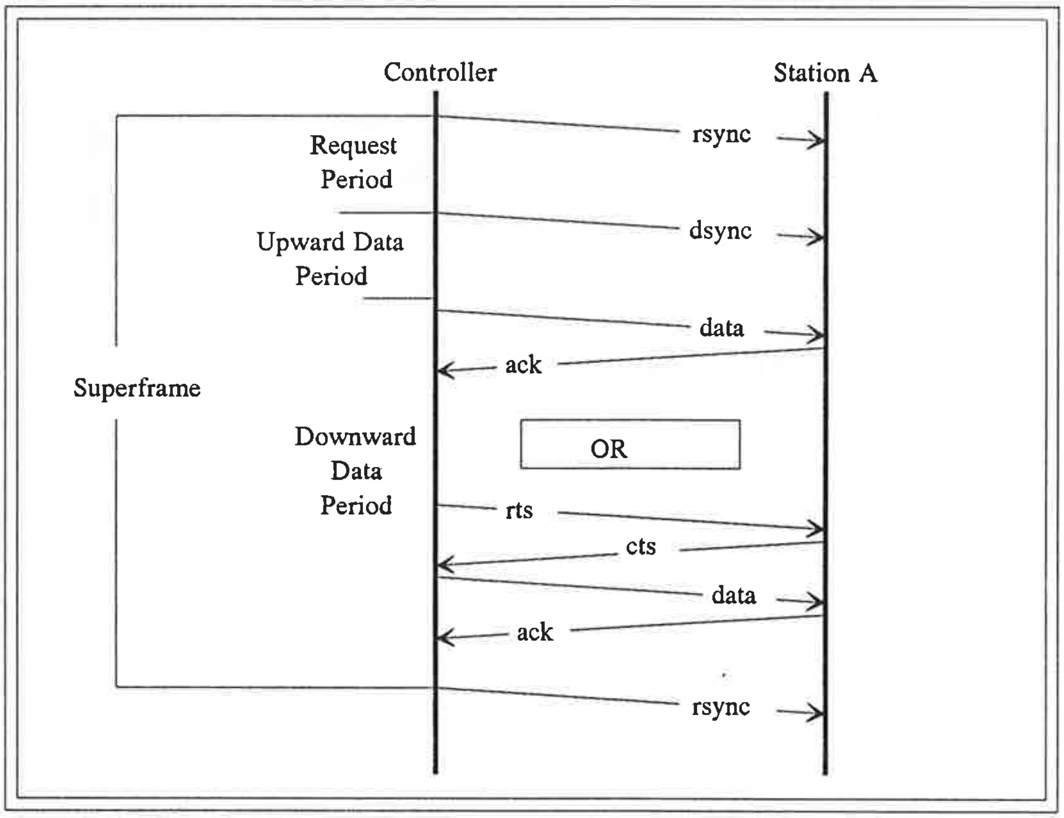

Figure 7.4 - Centralized Mode: Data from Controller to Station A

Any station that issued a request, during the previous request period, to send data to another station is allocated bandwidth by the controller to do so in this period. The controller sends a CTS to the requesting station indicating the ID of the destination station specified in the request. The requesting station may then send data to the destination station.

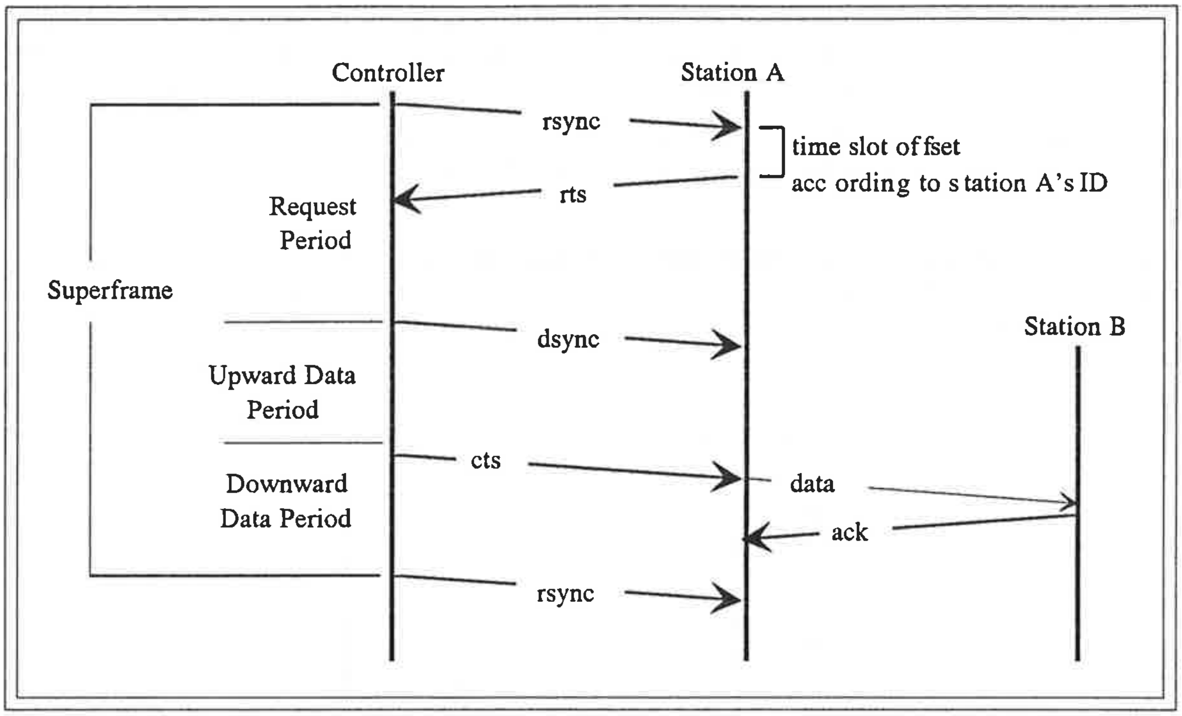

Figure 7. 5 - Centralized Mode: Data from Station A to Station B

7. Default Mode and Changing Modes

Distributed mode is the default operating mode of all non-controller stations. In distributed mode stations always listen prior to accessing the medium, for at least long enough to get an accurate Net Allocation Vector (NAV) built, so they will receive any synchronization frames which a controller within their range is issuing. When a station determines it is in a controller's BSA, it switches to centralized mode operation, the first step of which is to register with the controller.

There are two types of controllers, which have different default operating modes:

1. Dedicated controllers - stations with operating mode always centralized;
2. Potential controllers - stations with default operating mode distributed, which are capable of becoming centralized mode controllers.

The motivation for a potential controller to switch from distributed mode to becoming a controller is an implementation decision. One anticipated use of a potential controller is an AP which functions in distributed mode until it receives a Request To Send Time-bounded (RTST) frame. At that point this AP becomes a controller because centralized mode is required to support time-bounded service. Another might be a smart distributed mode AP which detects too high a rate of collisions due to a growing population in it's BSA, and switches to centralized mode to attempt to alleviate the situation.

8. Positive ACK and Duplicate Detection

Broadcast and multicast frames are not acknowledged, and never retransmitted. With that exception ...

Like the WHAT protocol, the Spectrix is a positive ACK protocol, in both operating modes. When a DATA frame is sent the source station expects to receive an ACK immediately following transmission. If an ACK is not received after a known period of time the data is scheduled for retransmission. After N tries the data is discarded and a failure reported.

The goals of the retransmission mechanism are: (1) to minimize retransmission by not re-sending the data when only the ACK was lost; (2) filtering out as many duplicates as possible; and (3) guaranteeing that no data will be mistakenly discarded as duplicate. This mechanism does not guarantee that no duplicates will be passed on to the client protocol.

Retransmission and duplicate detection are accomplished using three single bit flags in the control field of frames. These flags are kept by each station for each station with which they have been in contact, i.e. they describe the point-to-point transfer between two stations:

- The retry bit is zero on initial transmission and one on retransmissions. Note that re-sending an RTS due to not receiving a CTS is not considered a retry - the DATA frame must have been sent for a retransmission to occur;
- The sequence bit is an alternation bit, which can be thought of as a single bit sequence number;
- The out-of-sequence bit invalidates the previous sequence bit. If this is the first data sent from station A to station B, the out-of-sequence bit is set. The out-of-sequence bit is also set after an expiration of the N retry count, because the sender no longer knows the last sequence bit received by the destination.

In the following diagrams, which illustrate the use of these three flags, each frame is followed by the state of the three flags. E.G. RTS(x,y,z) - in this notation x = retry flag; y = out-of-sequence flag; and z = sequence flag.

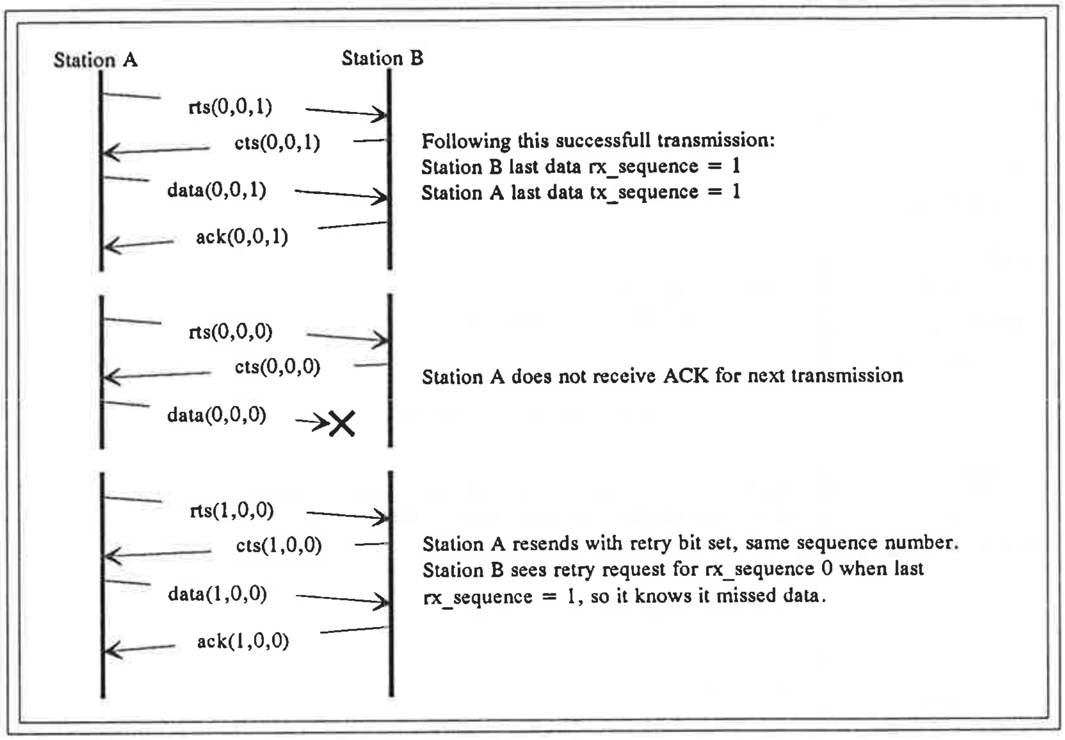

Figure 8.1 - Retransmission due to DATA frame loss

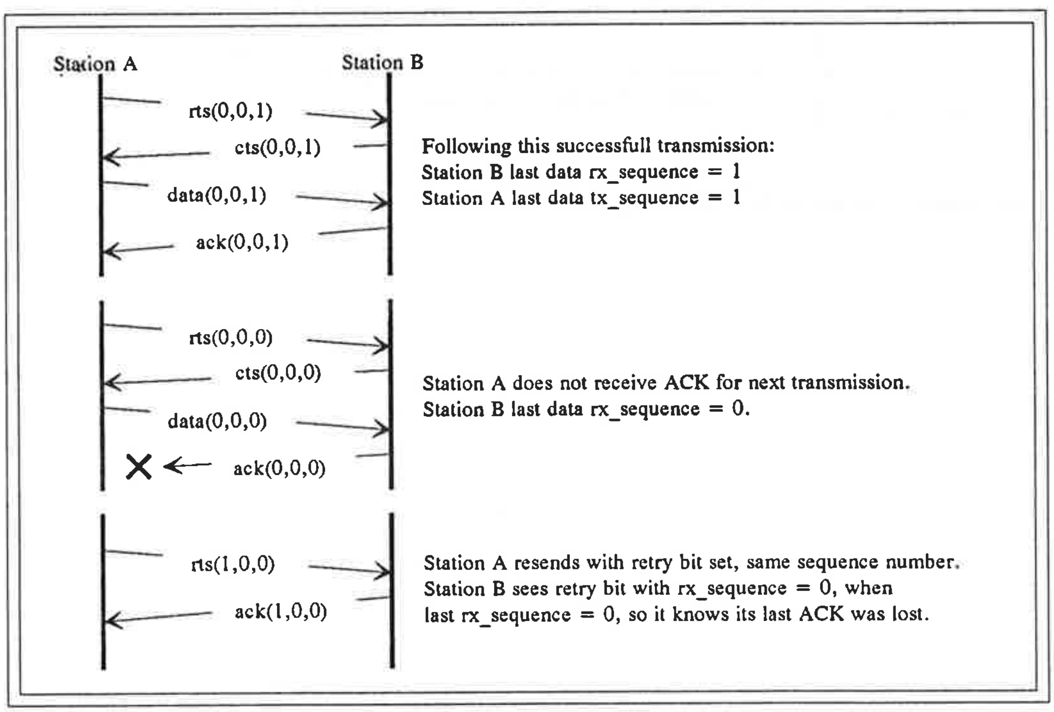

Figure 8.2 - Retransmission due to ACK frame loss

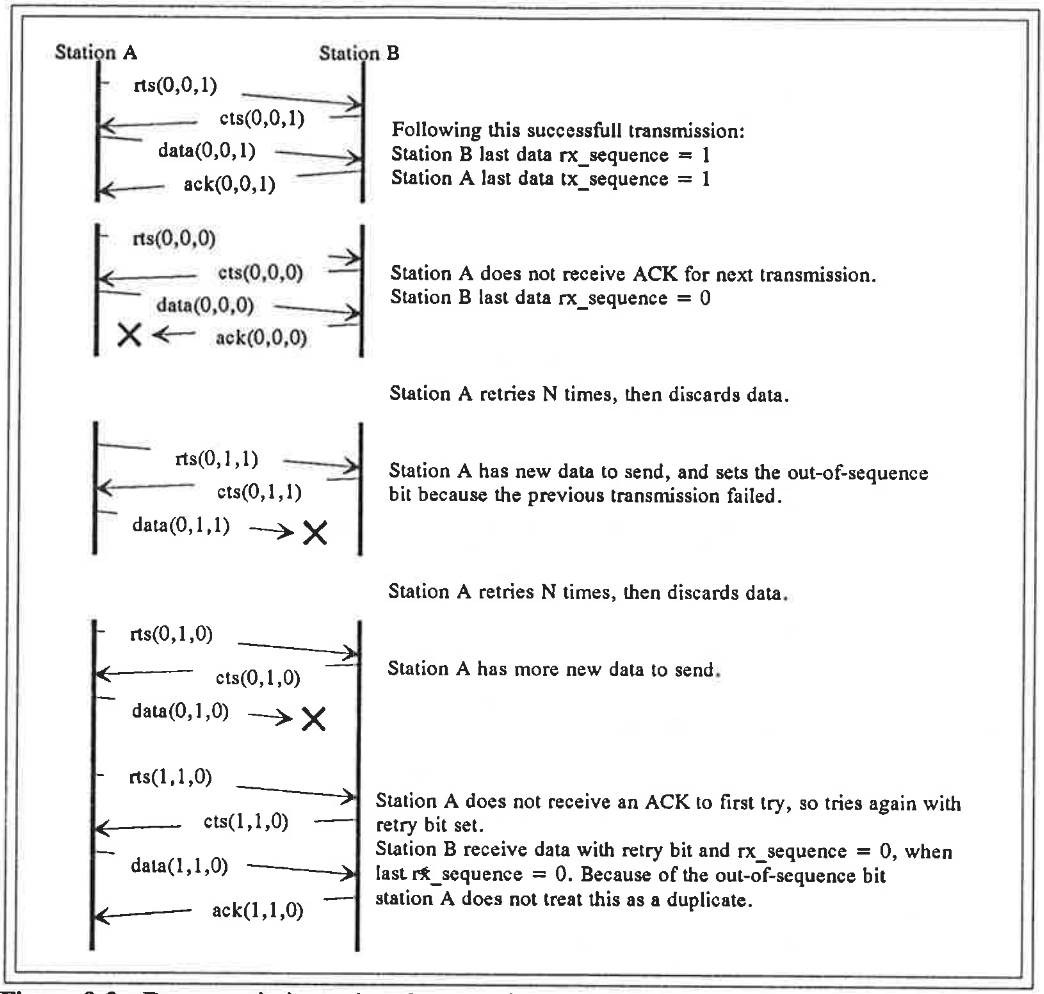

Figure 8.3 - Retransmission using the out-of-sequence bit.

9. Overlapping Operating Modes

9.1. Distributed Mode Station Overlapping with Centralized Mode Station

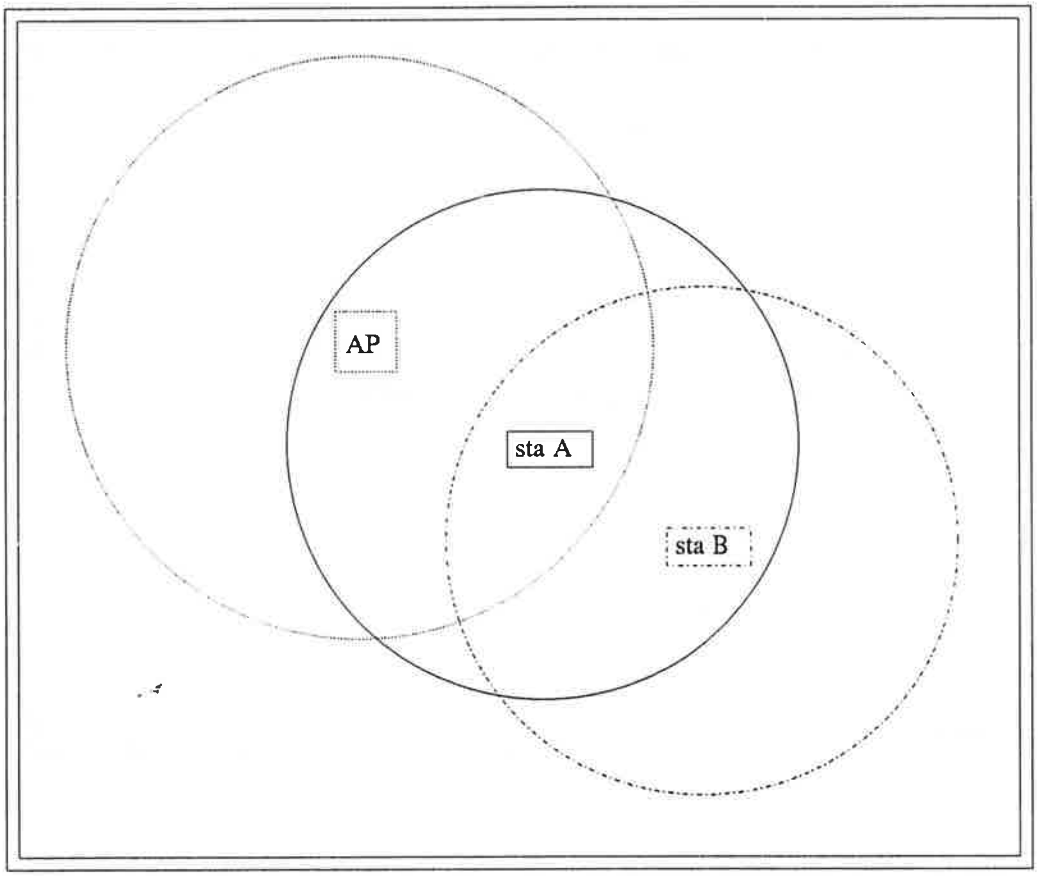

Figure 9.1

In the diagram above, assume that the AP is a controller station, so station A is operating in centralized mode. Station B is not within range of a controller, so it is operating in distributed mode.

Station B is operating in distributed mode, so it is listening to everything that comes from station A. It can tell from the destination ID of request frames from station A that station A is operating in centralized mode. What this means to station B is that when station A issues an RTS the associated CTS and DATA frames are not forthcoming immediately. It has no way to judge when they will come, so it has to consider the medium unavailable until they are seen.

Station B may see the medium as available during the request period, and if station A does not issue a request this is no problem. But if station A does issue a request it may collide with station B's frames, potentially interfering with station B's communication, but not affecting station A's request.

Station B may also see the medium as available during the AP's downward data period, and cause interference with the transfer of data from the AP to station A.

In summary, the performance of both the stations operating in distributed mode and the stations operating in centralized mode will be degraded by this type of overlap. However, although collisions may occur, all stations are still able to operate.

9.2. Overlapping Centralized Mode BSAs

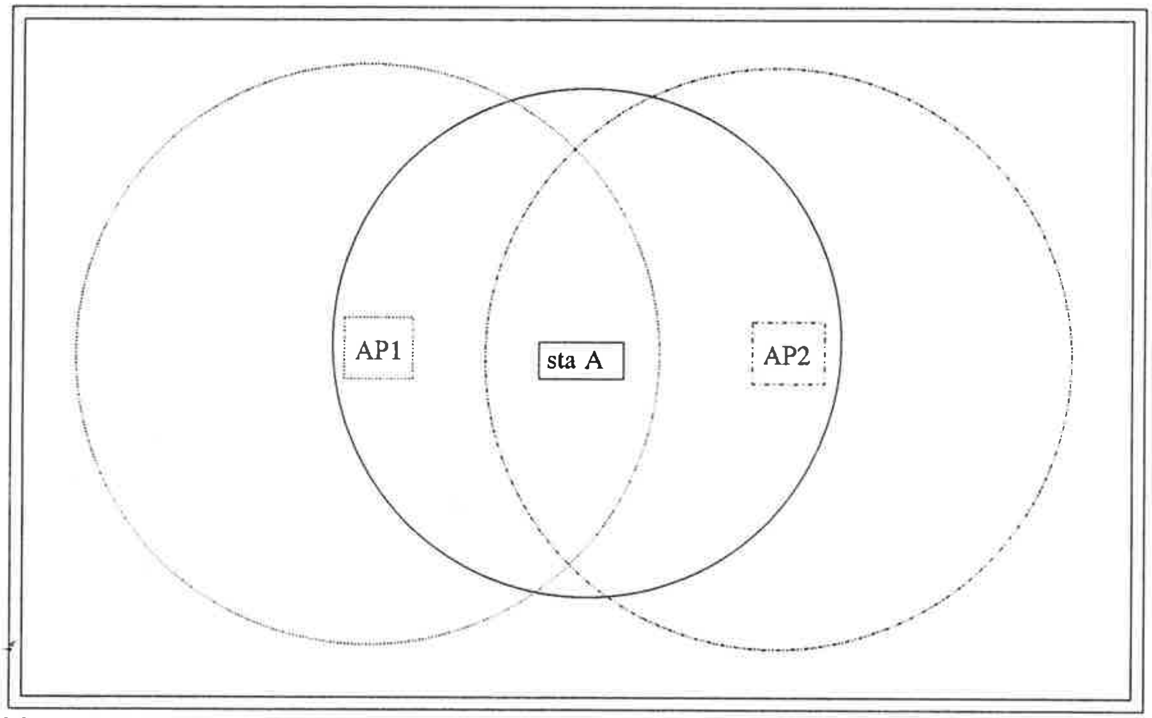

Figure 9.2

In the diagram above, assume that station A came alive, listened, and heard synchronization frames on the same channel from both AP1 which is operating as a centralized mode controller, and AP2 which is also operating as a centralized mode controller.

Station A cannot register with either AP without interfering with communication in the other AP's BSA. Station A could operate in distributed mode, but it will never find an opportunity to capture the media.

The proposed action in this case is for station A to register with one of the APs (causing repeated interference in the other until it successfully does so), with an indication in the registration request of the problem. Once one AP knows of the situation it can take some action. The station can also notify higher layers that it is not able to communicate.

Now assume that station A was already registered with AP1, and moved into range of AP2. This is even worse, because station A may never become aware of the situation. It has its receiver off during AP1's request period, and unless it has something to send, also off in the upward data period. In the downward data period it is most likely only listening to frames which are for it's own DID. So it may not see traffic from AP2's BSA. If it does, it can take the same action as described previously, alerting AP1 and its own higher layers of its inability to communicate. But until then, it is an interferer in AP2's BSA whenever it transmits in AP1's BSA.

10. Possible Enhancements

During the course of developing the centralized mode operation quite a number of suggestions/questions came up. Some yielded ideas which could be incorporated in the protocol, depending on how complicated one wants to make it.

1. "My number one concern is power consumption - I don't want to have to bring my receivers up to receive both the RSYNC and DSYNC, just for the downward data period."

Fields could be added to the RSYNC to accommodate this, or there could be different frame type called an expanded RSYNC (ERSYNC). The ERSYNC could specify the total time of the request period and upward data period, and stations receiving an ERSYNC would know not to expect a DSYNC.

The true length of the upward data period is not known at the time of ERSYNC transmission, but power consumption is considered so important that the implementer is willing to sacrifice bandwidth. The controller could always use a fixed length for the upward data period, or it could estimate the length based on previous traffic - this would be up to the implementation.

2. "My number one concern is power consumption - I don't want to bring my receivers up unless I know there is data for them."

Fields could be added to the DSYNC for this or an EDSYNC frame could be created. The EDSYNC frame could contain a destination map indicating which stations must keep their receivers on during the downward data period. This means that data received from the distribution system to go to stations after the map has been sent must wait for the next superframe - power consumption must be more important than transmission latency (and overhead, because the map could be large!) to implementers using this approach.

3. "I have high population density and/or time-bounded requirement so I want to use centralized mode. But I don't care about power consumption, so I don't want to waste overhead."

Use an ERSYNC, so you don't have to send the DSYNC. Specify in it that the length of the upward data period is zero, so all receivers will come on after the request period and stay on until the next RSYNC. Then you can mix your upward and downward data as you choose in the data period because all receivers are on.

Send your DATA frames from the controller without the preceding RTS/CTS, if your desire for saving bandwidth outweighs the possibility of having overlapping distributed mode stations which may need that information to minimize interference.

4. "I want to have AP/controllers communicating over the distribution system so that I can overlap centralized mode BSAs using time isolation. But to save power I want to turn off the receivers in one BSA during the superframe of another."

Add another type of synchronization frame, a pause sync (PSYNC) which specifies a time length for which receivers of stations registered with the controller sending it can turn off.

5. "It may take more than one IR transceiver (let's say 2) to cover a large room, but we want the room to be one centralized mode BSA. We can't run the two controller's transceivers simultaneously due to multipath interference."

Change the structure of the superframe a little. It is still composed of periods delimited by synchronization frames, but do two RSYNCs and two DSYNCs, each containing the same superframe number. Send an RSYNC from the first transceiver, get the request list from it. Send an RSYNC from the second transceiver, get the request list from it. Use the quality-of-signal information associated with each request to determine which transceiver is better for communicating with which station. Then do a DSYNC from one transceiver and service the stations that have better quality from it, then a DSYNC from the other and service the other stations.

This method has high overhead, because the request period was done twice. The total data period is only longer by one extra DSYNC.

This leads to the conclusion that the superframe can be composed of as many request periods and data periods as desired. The sync frames should contain a superframe number, so that stations know when to retransmit because they didn't get serviced in this superframe.

6. "I expect all my stations to register at the beginning of the day, and none, or very few after that."

When the controller has few stations registered, have a lot of registration slots available. When it has a lot of stations registered, have few or have only every third or fourth request period have any registration slots in it. Or when there are a lot of stations registered have complete overlap between owned and registration slots.

7. "I want the controller to able to know where all stations are at any given time, whether they have communicated lately or not."

Add a mandatory-response RSYNC frame (MRSYNC) and send it out periodically. When a controller issues a MRSYNC frame all registered stations must respond in their owned request slot. If they had not intended to generate a frame in their owned slot they must send a forfeit frame (FORF) there which has no effect other than to note their presence.

This method could also be used to age time slot ownership.

8. "All registered stations own a time slot in the request period, couldn't they send their ACK frame there in the first request period after they receive data from a controller?"

YES - the ACK could be sent at that time rather than immediately following data, and there could be a combined request/ack frame (or a control flag in the request frame). This would definitely save overhead.

BUT - without this enhancement the action of a station on receipt of a DATA frame is always the same (send an ACK frame immediately), regardless of centralized or distributed mode, and regardless of the DATA frame source.

11. Basic Frames

Request Sync

Preamble	
SD	
DID	⇐Broadcast
Type	⇐RSYNC
Control	⇐Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐Station Identifier, ID of originating controller station
TotalSlots	⇐Total number of time slots in SyncPeriod (including RegSlots)
RegSlots	⇐Total number of time slots which are for registration only
SuperFrame	⇐Superframe number
FCS	
ED	

Data Sync

Preamble	
SD	
DID	⇐Broadcast
Type	⇐DSYNC
Control	⇐Control flags: AP, sequence, out-of-sequence, retry, hierarchical
UpLength	⇐length of Upward Data Period
SuperFrame	⇐Superframe number
FCS	
ED	

Request To Send,
Request To Send Indirect,
Request To Send Time-bounded

Preamble	
SD	
DID	⇐Destination Station ID
Type	⇐RTS, RTSI, or RTST
Control	⇐Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐Source Station Identifier
DataLength	⇐Length, in octets, of data the source wants to send
DA	⇐Address station to which data is to be sent, 48-bit address
FCS	
ED	

Clear To Send

Preamble	
SD	
DID	⇐ Destination Station ID
Type	⇐ CTS
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
CTSSID	⇐ Clear to Send to Station Identifier
DataLength	⇐ Length, in octets, of the data the destination station is to send
FCS	
ED	

Registration Request
Registration (Time-bounded) Request

Preamble	
SD	
DID	⇐ Destination Station ID
Type	⇐ RREG or RTREG
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐ Source Station ID (registration time slot number)
SA	⇐ Address station registering, 48-bit address
FCS	
ED	

MPDU Data

Preamble	
SD	
DID	⇐ Destination Station ID
Type	⇐ DATA
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SA	⇐ Address of data originator, 48-bit address
DataLength	⇐ Length, in octets, of data to be sent
Data	⇐ Data
FCS	
ED	

Acknowledge

Preamble	
SD	
DID	⇐ Destination Station ID
Type	⇐ ACK
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐ Source Station Identifier
FCS	
ED	

Mac Management Data

Preamble	
SD	
DID	⇐ Destination Station ID
Type	⇐ MDATA
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SA	⇐ Address of data originator, 48-bit address
MType	⇐ Type of MAC management message
Data	⇐ according to MType
FCS	
ED	

MType	Data
Registration Accept	SID, SA
Registration Reject	none
Registration Cancel	SID, SA

12. Possible Enhancement Frames

Pause Sync

Preamble	
SD	
DID	⇐ Broadcast
Type	⇐ EDSYNC
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐ Station Identifier, ID of originating controller station
PauseLength	⇐ length of receiver off period
FCS	
ED	

Extended Data Sync

Preamble	
SD	
DID	⇐ Broadcast
Type	⇐ EDSYNC
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐ Station Identifier, ID of originating controller station
UpLength	⇐ length of Upward Data Period
DownMap	⇐ Map of stations to receive data in downward data period
SuperFrame	⇐ Superframe number
FCS	
ED	

Extended Request Sync

Preamble	
SD	
DID	⇐ Broadcast
Type	⇐ EDSYNC
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐ Station Identifier, ID of originating controller station
TotalSlots	⇐ Total number of time slots in SyncPeriod (including RegSlots)
RegSlots	⇐ Total number of time slots which are for registration only
UpLength	⇐ length of Upward Data Period
DownMap	⇐ Map of stations to receive data in downward data period
SuperFrame	⇐ Superframe number
FCS	
ED	

Mandatory-response Request Sync

Preamble	
SD	
DID	⇐ Broadcast
Type	⇐ MRSYNC
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐ Station Identifier, ID of originating controller station
TotalSlots	⇐ Total number of time slots in SyncPeriod (including RegSlots)
RegSlots	⇐ Total number of time slots which are for registration only
SuperFrame	⇐ Superframe number
FCS	
ED	

Forfeit

Preamble	
SD	
DID	⇐ Destination Station ID
Type	⇐ FORF
Control	⇐ Control flags: AP, sequence, out-of-sequence, retry, hierarchical
SID	⇐ Source Station Identifier
FCS	
ED	

Appendix A - An Implementation of Spectrix' Reservation/Polling Protocol

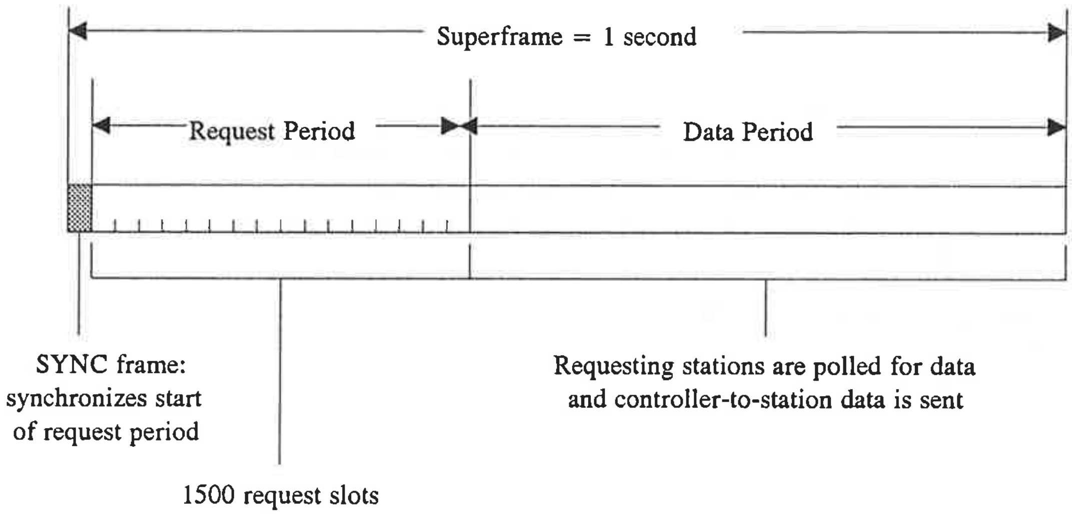

Stations have fixed addresses, 0 to 1499. A station's address is the number of its owned time slot. When a station has data to send it issues a request in its owned time slot. In the data period the controller polls all stations which have issued requests, and sends data to stations.

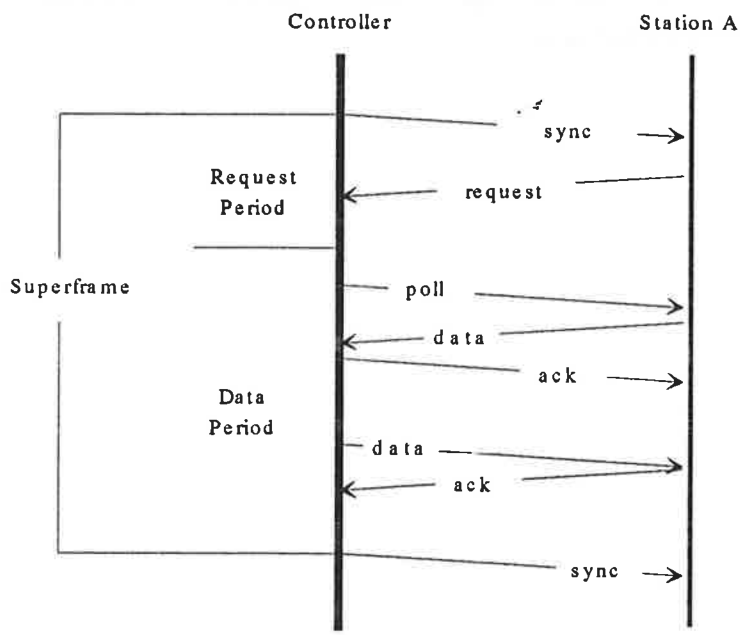

- Direct station-to-station communications not supported.
- Controllers are Access Points connected to existing wired LANs.
- Infrared PHY, single channel. Network isolation by space - BSAs do not overlap.

Appendix B - 802.11 Issues (21 criteria item # in parenthesis)

3.1 What is the impact of the MAC implementation complexity in regard to time-to-market? (15)

The CODIAC protocol can be implemented in many levels of complexity. Where time-to-market is of primary concern a simple implementation could be chosen to accomplish this.

4.5 Can a station be a member of an ad-hoc and non-ad-hoc network at the same time?

Yes. Station A can be registered with a controller/AP, and associated with that AP - a member of an infrastructure network. Station B may be registered with that controller/AP, but not associated with the AP, it is registered only for the purpose of conversing with other wireless stations - it is not a member of the infrastructure network. These two stations can converse without station A having to dis-associate from the AP, so it retains its membership in the infrastructure network while forming an ad-hoc network with station B.

5.3B What logical functions are needed to support the defined infrastructure services?

These services are defined in closed Issue 5.3A as: association, re-association, disassociation, authentication, privacy, integration, and network management.

For any of these services which require exchange of information over the wireless medium, the CODIAC protocol proposes using MDATA frames. Because delivery of these frames is critical, they are transferred in the four-step transaction in the same manner as client data. These frame formats are yet to be fully defined.

Association, re-association, disassociation, and integration all require an AP. These services are supported by the AP bit which is set in frames sent by the AP, which also serves to notify stations of its presence.

5.9 How to determine that Access Points (APs) are present?

All frames are marked with an AP bit which indicates that they originate with an AP. If a station listens and does not hear frames from an AP, it can send a broadcast RTS with the Hierarchical bit set, which indicates that the RTS is intended for an AP only - this will cause any AP present to identify itself.

6.3 Unauthorized access impact MAC throughput. (1)

The philosophy of the CODIAC protocol centralized mode is that registration with a controller is completely separate from association and authentication with an AP. The purpose of registration is to facilitate coordinated sharing of the bandwidth - stations cannot be precluded from this on the basis of authorization. The example of the centralized network that covers the park across the street from the building in which it is implemented is such a case. Stations in that park must register with the controller to share the bandwidth - they must not be excluded from using that bandwidth.

In distributed mode the bandwidth is available to all stations, the registration process does not exist.

In either mode stations get associated with and authorized to use the infrastructure through an AP by an association/authentication process. Unauthorized stations can repeatedly attempt to get authorization and be rejected, and this will impact throughput by adding a repeated transaction.

9.2 What are the area coverage implications of MAC timing constraints? (10)

On the assumption that this issue arose from the Ethernet maximum cable length specification which is driven by the timing constraints of CSMA/CD: No timing constraints are imposed by this protocol that would limit coverage area of LAN dimensions.

9.3 Must the same MAC work in a minimum system and maximum system (network size independence)? (16)

Yes. Not just to work in both, but to work efficiently in both is the goal of the CODIAC protocol.

10.1 What Coordination Function (CF) will be specified in the standard?
10.2-B Do multiple CFs need to be specified?

Both Distributed Coordination Function (DCF) and Point Coordination Function (PCF) are required to support efficient operation with network size independence for asynchronous service. PCF is required for TBS, but this should not be forced on small population and ad-hoc networks.

10.2-A What are the events that cause switching between multiple CFs?

1. Switch from DCF to PCF: Request for Time-bounded service from a station to a controller which supports TBS.

2. Possible implementation, switch from DCF to PCF - detection of high traffic causing high rate of collisions.

10.3 What are the issues surrounding the Point Coordination Functions (PCF) and Distributed Coordination Function (DCF) arguments?

1. PCF is required for TBS support.

2. DCF facilitates ad-hoc networks better because it does not require a controller.

3. PCF is better than DCF for minimizing power consumption of portable stations.

4. PCF is better for high population networks, deterministic media access to avoid collisions.

5. DCF is lower overhead and possibly lower access delay (in small population BSAs).

11.3 Is there a need for multiple APs per Basic Service Set (BSS)?

Although no need is envisioned, no reason for preclusion is seen. With the CODIAC protocol only one controller per centralized mode BSA is required, but any number of stations could be APs.

13.3 What support will the standard provide for power management:
13.3A DC power (power consumption)?
13.3B RF power (signal strength)?
13.6 How will the MAC standard address Power Consumption? (9)

13.6 & 13.3A Power Consumption

Some implementations are more concerned with power consumption than others. The CODIAC protocol allows implementations to trade off power consumption requirements with overhead and access delay. These features are described in the main text of this document.

13.7 & 13.3B Signal Strength

Section 10, point 5 addresses one way in which the centralized mode may be used to aid in signal strength management. No investigation has been done in this area, research and development may uncover more.

14.4 Ability to establish peer-to-peer connectivity without prior connection (e.g. without "knowledge of the presence of your peers"). (2)

Interpretation - can a station initiate communications with another station without knowing that it is present, and what its wireless address is?

Yes. In the RTS frame contains the 48-bit address of the intended destination station. In distributed mode this frame is broadcast, so the destination station can respond if it is there. In centralized mode the RTS is sent to the controller, and it can use its knowledge of registered stations to determine the wireless address of the destination.

Also, use of the AP bit and the Hierarchical bit allow stations to identify APs without any prior knowledge.

15.6 What is the algorithm for managing partitioning of capacity between Time-bounded and Asynchronous services?

That should be left to the discretion of the implementation. The CODIAC protocol allows different implementations to tailor servicing of stations to their needs while still remaining compatible.

15.8 Do all stations and all infrastructures support the Time-bounded service?

a) Stations

The CODIAC protocol requires that all non-controller stations be well behaved in both operating modes. This means a station must be: (1) capable of communicating in both modes; or (2) capable of communicating by the distributed mode rules only, but it must be quiet in the presence of a controller; or (3) capable of communicating by the centralized mode rules only, but it knows it must be quiet when it does not hear a controller.

This means that for non-controller stations "supporting" (where "supporting" means not precluding other stations from using TBS) TBS with the CODIAC protocol is a given, because TBS is provided by centralized mode operation .

For controller stations, whether they can operate in both modes should be an implementation decision. However if a station requests TBS, there should be a specific negative response to that request if the service cannot be provided (not yet defined).

b) Infrastructures

Yes, where the definition of support is to handle in a well behaved manner - i.e. where a station requests TBS there should be a negative response to that request if the service is not provided.

If support = provide, then No.

Summary - in agreement with Pro arguments 3.1 and 3.5

15.9 How will the standard address the MAC ability to service various traffic: data, voice, and video. (6)

The CODIAC protocol supports asynchronous and time-bounded services. The centralized mode can be implemented to support the requirements of various TBS time constraints.

17.2 What level of reliability for Broadcast (Multicast) Addressing is required? (20)

Multicast and broadcast reliability is directly tied to the MSDU error rate, as they cannot be acknowledged. This is the case for all LANs, wired and wireless. These are not inherently reliable delivery mechanisms.

17.3 What is the extent of Multicast? (BSS or ESS)

Both ESS and BSS multicast should be supported, a station should be able to explicitly control the scope of multicast (this supports the position of document 93/40 on the WHAT protocol). The hierarchical bit provides this capability.

17.5 What is meant by address: size? is IEEE 802 addressing OK?

IEEE 802 addressing is required (supports the position of document 93/40 on the WHAT protocol). Wireless stations should be identified by 48 bit unique IDs that are compatible with other IEEE 802 standards. The 48 bit addresses of source and destination stations are contained in the four step transaction of the CODIAC protocol.

18.3 Will the standard support PHY with variable rates?

Yes. RSYNC frames could be issued at different rates within a superframe, or different superframes could be issued. PSYNC could be issued at one rate while communication was going on at another.

Little consideration has been given to this issue at this time. However, this is a very important issue. First generation wireless LANs will be released at lower speeds than forthcoming generations, but they must coexist - it is not desirable tell customers they must upgrade their equipment because the company across the hall installed a newer, higher speed LAN.

19.1 Shall the 802.11 standard depend on the layers above the MAC for recovery from failed transmits? If so to what extent?

Partially. A retry mechanism should be implemented in the MAC as required to bring the MSDU loss rate up to the equivalent of wired LANs. (See Issue 19.5)

19.2A Will the IEEE 802.11 MAC look like all other 802 MACs regarding delivery reliability?
19.2B How does Multicast affect this decision?

19.2A - Yes - see 19.5

19.2B - Broadcast and Multicast will not be as reliable - see Issue 17.2

19.5 What kind of error recovery mechanisms are to be incorporated into the MAC?

Supports the position of document 93/40 on the WHAT protocol - the 802.11 MAC should include a positive acknowledgment protocol with low level retries. This mechanism helps the MAC present approximately the same level of MSDU delivery reliability as other IEEE 802 protocols.

19.6 What is the strategy for capacity control?

The CODIAC protocol is in itself a strategy for capacity control. The purpose of the two operating modes is allow efficient media use under different capacities, and in centralized mode each implementation's strategy for management of request periods and data periods in centralized mode is its strategy for capacity control.

19.7 Is the maximum number of stations to be specified? If so how many? (5)

No. That should be up to the implementation.

In distributed mode the protocol will begin to break down at a certain number of stations, and the implementer should decide what action to take about that - whether to switch operating modes, or to make the degradation limit a parameter of the network.

In centralized mode it is a function of the intended application. An application with huge numbers of stations with small payload and/or tolerance for large transfer delays can be supported, as can an application with smaller population with need of shorter transfer delays. The CODIAC protocol can be set up to accommodate either, without loosing compatibility.

19.8 How will the standard address the MAC robustness in the presence of co-site dissimilar networks? (8)

On the assumption that "dissimilar" means not so different that they don't see each other (e.g. IR and SS), and not so similar as to be able to recognize each other's MSDUs - Co-site dissimilar networks interfere with each other. There is nothing the MAC can do about this that is different from handling interference of any other kind.

19.10 How will stability under heavy load be addressed?

The centralized mode of the CODIAC protocol remains stable under heavy load by increasing transfer delay. This is further explored in document "Performance of the CODIAC protocol".

19.11 How will transmission lost be addressed?

Issues 19.1 and 19.5 cover this issue.

The CODIAC protocol proposes positive ACK and retransmission to bring the transmission loss rate to approximately the same level of MSDU delivery reliability as other IEEE 802 protocols.

20.2 Can MAC handle different preamble lengths from different PHYs?

No. If different PHYs must generate different preamble lengths then preamble should be handled by the Medium Independent Layer, which is on the PHY side of the MAC/PHY interface. The preamble would be stripped off by the time the frame is seen by the MAC.

To facilitate MAC independence from preamble length, perhaps the preamble should not be considered part of the MAC frame.

20.3 What is the MAC frame structure?

The frame structure is designed with the following goals: (1) to minimize the frame size while keeping a consistent frame structure; (2) to have a minimum size destination identifier at the start of the frame to allow destination determination of frames as quickly as possible; (3) to provide a level of error detection suitable to the high bit error rate of the wireless media.

Minimum Frame Length (12 +n) octets								
		FCS Coverage						
Preamble	SD	DID	Type	Control	INFO	FCS	ED	⇐ field name
8n	8	16	8	8	(optional) 8m	32	8	⇐ field length (bits)

Preamble = Preamble (n to be determined)
SD = Start Delimiter
DID = Destination Identifier
Type = Frame Type
Control = Control Flags: AP, sequence, out-of-sequence, retry, hierarchical
INFO = Information (0 < = m < = to be determined)
FCS = Frame Check Sequence, CRC-32
ED = End Delimiter

20.4 How is the MAC time preservation ordering of SDU to end systems (LLC requirement) addressed by the standard? (21)

The CODIAC protocol is a stop-and-wait ARQ, it does not change the order of MSDUs.

21.2 How will the MAC robustness in the presence of non-reciprocal wireless medium be addressed in the standard? (14)

If this means stations may have different receive and transmit coverage area:
In CODIAC protocol centralized mode, if the relationship between the controller and a station is asymmetric the station will not be able to register. Minimal bandwidth will be lost as it repeatedly tries to do so. In distributed mode the RTS/CTS exchange will fail, avoiding the wasted bandwidth of attempting to send the data itself.

If this means non-reciprocal traffic load:
The CODIAC protocol is flexible in the assignment and duration of the data periods in centralized mode, both at run-time and per implementation, creating no problems handling non-reciprocal traffic loads. This is a moot point for distributed mode as it has no directionality.

25.1 Will the standard provide a procedure to reserve medium channel capacity?

Not a lot of work has been done so far in this area, however this facility can easily be incorporated into the CODIAC protocol by adding information to the request frame specifying a reservation of a particular length, or even making a "connection request" for a certain amount of bandwidth which could stand as a reservation of channel capacity until the connection is torn down, rather than having to issue a request every superframe.

25.2A Must the MAC work on a single channel PHY?
25.2B Will the standard support multiple channel PHYs?

Yes and yes.

25.5 What is the definition of MAC fairness of access? (11)

The definition of fairness of access is all stations having an equal opportunity to access the media. Things about a MAC that can make access opportunity unfair are:

(1) sensitivity to the near/far bias (capture effects);
(2) allowing one station to hold the medium once it has it;
(3) bias to a particular data path - AP to station; AP from station; or station to station;
(4) bias to a traffic type, TBS or asynchronous.

The CODIAC protocol addresses these items:

(1) see Issue 25.6.
(2) Maximum frame length controls this to in both modes. In distributed mode once a station has made a transaction, of up to maximum length, it must re-contend for the medium like all the other stations. In centralized mode the controller implementation controls this fairness. At the end of the request period it has the information required to divide up the data period bandwidth as it sees fit.
(3) In distributed mode there is no distinction between these data paths. In centralized mode the controller implementation controls this.
(4) In both modes the AP implementation controls this. An AP could deny a TBS request if it feels that the asynchronous traffic is being unfairly denied access by the amount of TBS traffic.

25.6 How will the standard address the MAC facilitation of 'access fairness' (insensitivity to near/far bias)? (12)

In the CODIAC protocol centralized mode sensitivity to the near/far bias will only come into play in the registration slots. If two stations attempt to register in the same slot and one of them has signal strength enough to obliterate the other, the winner will get registered and the loser will have to try again next superframe.

Summary - (1) the near/far bias can cause a minor delay in registration, but the protocol is insensitive to it for data transfer in centralized mode; (2) Distributed mode is sensitive to the near/far bias during the RTS/CTS exchange.

26.1A Does the concept of priority need to be addressed in the MAC?
26.1B Different traffic priorities?
26.1C What is priority? (13)

26.1C - Priority is a station having better access to the medium, in terms of access delay and/or time length of access, than other stations.

26.1A - If the concept of priority is addressed in the MAC: The CODIAC protocol lends itself very well to the implementation of priority in centralized mode. If priority is added to the RTS frame then the controller can service requests in prioritized sequence in the data period. The controller can also assign quantity of bandwidth to requesting stations in a prioritized fashion. Priority is not a concept which can be applied to the CODIAC protocol distributed mode.

26.1B - With respect to traffic types, in distributed mode TBS traffic is not supported so it is not relevant. In centralized mode the protocol does not give priority to either traffic type, but an implementation could do so, as TBS requests are marked.