IEEE P802.11
802 LAN Access Method for Wireless Physical Medium

DATE: August 23, 1991 **REVISED:** September 4, 1991

TITLE: **SEQUENTIALLY-USED COMMON CHANNEL ACCESS METHOD**

AUTHOR: Chandos A. Rypinski,
Chief Technical Officer
LACE, Inc. Telephone: 707 765 9627
921 Transport Way Facsimile: 707 762 5328
Petaluma, California 94954 USA

SUMMARY

This access method uses a single high-data-rate channel sequentially at overlapping and interfering Access-points (one reuse group) all under the control of a single Access Manager within a multi-port Hub Controller. In general, the capacity can be equal to the parallel use of a lower rate at the same number of Access-points in a channelized system.

The primary plan requires infrastructure, but permits direct Station-to-Station transfers when it is possible. The plan provides connection-type service consistent with IEEE 802.6 cell format.

A secondary version of the plan is given for operation without infrastructure and without support of connection-type services.

The previous contribution (IEEE 802.11/91-19) on access method is superseded in detail use of messages by the current contribution (IEEE 802.11/91-95), however, the original contribution contains material on calculated efficiency and compatibility with 802.6 ATM cell transfer which remains relevant though slightly inaccurate because of increases in message lengths to allow channelization and multiple simultaneous connections at one Station. The purpose of this revision is to present the single time-shared channel concept consistently with channelized systems elsewhere presented, and to incorporate the common message set for all of the methods.

This method is the first of three access methods that have been developed all using the same message set and which can be characterized as follows:
1) **With and without infrastructure, sequential use of one channel at all Access-points within one reuse group.**
2) With infrastructure control, sequential use of a common setup channel and parallel use of a number of data transfer channels derived by code-division spread spectrum or otherwise.
3) Independent of infrastructure and without virtual circuit support, random contention use of a common setup channel and distributed channel selection for following use of one of several data transfer channels.

Table of Contents   Page

SEQUENTIALLY-USED COMMON CHANNEL ACCESS METHOD

NOTICE: The following paper is a revision, rewrite and abridgment of IEEE 802.9/91-19 "Access Protocol for IVD Wireless LAN." The purpose is to present the single time-shared channel concept consistently with channelized systems elsewhere presented, to incorporate a common message set for all of the methods, and to add a non-infrastructure secondary access method.

OVERVIEW

Three access methods have been developed all using the same message set.

This plan is dependent on infrastructure; however, the last section of this plan adds a non-infrastructure access method using the same message set.

The protocol is implemented with sequential message transfers using a message set which is capable of serving this and the other access methods.

Full provisions have been made for virtual circuits with demand-assigned bandwidth.

The meaning of "asynchronous" as used is that there are no fixed dimension timing frames or slots. Each transmission begins when the necessary pre-conditions are satisfied none of which are fixed time intervals. There is no addressing in the form of time location in the bit stream.

The topology assumed is one system (reuse group) with 9, 16 or a small arbitrary number of Access-points on which Invitation-to-transmit appears sequentially.

The protocol is fully described, but traffic capacity analysis is left for another contribution.

This access method is applicable to medium signaling rates of 8-20 Mb/s, and the description assumes 10 Mb/s.

FUNCTIONAL OBJECTIVES

The scope and intentions of this access protocol are as follows:
1) to operate equally on wire, radio or optical mediums.
2) to be suitable for indefinitely large systems carpeting a large area, or to be usable as an independent small system with many like type systems on all sides.
3) to limit energy density radiated when all of the simultaneous radiation is summed for interference effect on other spectrum users.
4) to attain full use of channel time under peak load conditions.
5) to avoid Station logic dependent on system configuration and public network interface protocols.
6) to make all Station logics identical independent of the size, scale or traffic capacity of the system in which they are used.
7) to make this local distribution system compatible with common backbone network technology.
8) To use the medium efficiently when the outward-inward traffic volume is asymmetrical,
9) To provide direct peer-to-peer communication when Stations are in range of each other.
10) To use no air-time to pad messages out to predetermined boundaries.

DESCRIPTION OF SYSTEM FEATURES

Architectural and logic choices give this protocol and the environment that it creates a special set of properties now described.

This access protocol is different from other 802 protocols in a number of ways mostly as a result of fitting to the particular conditions of use. The extent of an access group is smaller, and the probability of one packet being correctly received on the first try is much smaller. The repetition mechanism must be in the physical layer to bring the performance up to the accuracy levels of other 802 physical layers. By integrating acknowledgement and the repetition process into the access algorithm, the amount of time required for the execution of this function is greatly decreased when it is used.

Central Access-management and Asynchronous Access

By using a centrally-managed access method, use of air-time can approach 100%.

The right time to start a second use of the shared medium is immediately after a first use is concluded--and that is the purpose of asynchronous logic. Among shared mediums, radio is particularly diminished by unnecessary "dead" time.

A further property of radio is that there is inevitable interaction between contiguous illuminators. It is not possible to use one radio Access-point independently of those around it. This kind of logic is easily handled in common equipment, and unlikely to be effective or efficient distributed over the stations. Only at a common point is it possible to assemble and use essential facts about the system:

1) which Stations are available, and
2) capable of directly communicating with each other.
3) what outside priority traffic is waiting for transmission that could take precedence.
4) what is the local address corresponding to a global address in either 6 octet LAN or E.164 BCD format, and vice versa.

The entity providing central control functions is an "access manager" within a Hub Controller.

Central Control and Addressing

Stations have a long global LAN address, a short local address and may have a long E.164 telephone address. A directory function in the Hub Controller enables traffic within the system to use short addressing. Nonetheless, long addressing for both DA and SA will be used for externally addressed traffic.

When the *REQUEST* from a Station contains a long destination address and there is a *GRANT*, that packet or segment is completely defined by the short source address thereafter. This procedure significantly reduces overhead.

Short addresses are non-permanent and assigned automatically by the Hub Controller. They need not be used or apparent to Users.

Traffic originating outside the system will be received and forwarded with long DA and SA.

Automatic Repetition (ARQ) of Errored Packets

The advantages of ARQ are that it introduces no decrease in throughput except when it is needed, and it better compensates the temporary outage characteristic of the radio medium than would a full-time FEC.

Error-free (determined by CRC) packets originating at Stations are ACKnowledged by receiving Access-points. There is an automatic resend of the packet/segment limited to three attempts (two retries).

Virtual circuit packets may be allowed only one retry or none at all, and then dropped.

Stations are permitted to send *NACK* if they are addressed but the packet fails CRC.

Access-point Address and Moving Stations

Each Access-point is assigned a 4-bit identifier (API) which is broadcast in each *INVITATION-TO-REQUEST/REGISTER* message. (Note: 12 bits may be a better choice.) Each Station will know which Access-point is being received or used and whether the current

identifier is the same or different than the identified Access-point last used.

Signal-level Reporting From Access-points

The received signal level is reported back to the Hub Controller concurrently with reception. This signal level is used by the Hub Controller to decide which of several Access-points receiving the signal should be used for the next message exchange.

Asymmetry of REQUEST-GRANT Procedure

The inward *REQUEST-GRANT* procedure is essential to keeping Station logic simple and avoiding contention. It would be possible to use the same *REQUEST-GRANT* procedure in both directions rather than for only upward traffic.

802 philosophy is that a Station is always ready. A message is always sent--ready or not. It is for higher layers to provide a method of knowing whether Stations are there. To allow Stations to refuse a message, would be a point of incompatibility with 802 practice.

If Stations could say "wait," a further buffering and indeterminate delay factor would be introduced in the network implementation. The amount of storage required in this protocol is only that to do repeat send on no *ACK*. Otherwise messages "fall on the floor" after a multiple-try delivery attempt.

A symmetrical protocol is not used, because there is no functional value in downward *REQUEST-GRANT*, it would cost air-time, and it could create indeterminate buffer requirements in the infrastructure.

DESCRIPTION OF THE ACCESS PROTOCOL

This access method uses the message set shown in a separate contribution IEEE 802.11/91-80. (11AP16K) amended to add station-generated *GRANT*. This message set is also used or usable in other types of systems.

In this protocol direct peer-to-peer is a supported function with or without infrastructure.

The description below is from the viewpoint of the individual user Station unless otherwise noted.

Initial Conditions–
With and Without Infrastructure Present

If an infrastructure is present, its Access-points are transmitting *POLL*, *INVITATION-TO-REGISTER* and *INVITATION-TO-REQUEST* messages as described later below.

When a Station does not hear any of the above infrastructure messages, there are two possible alternative reactions:

1) the Station may assume a default no-infrastructure, contention mode is in operation, or

2) the Station activates a simplified default access manager which makes the first activated Station operate sufficiently like one Access-point to enable autonomous intercommunication within small groups of Stations.

Operation with no infrastructure is described in greater detail later below.

Registration Function

Stations just entering the system listen first for the *INVITATION-TO-REGISTER* (007) messages from which they identify the system providing the infrastructure. The Station cannot know which is the strongest signal received or which site would provide it. Though the Station can recognize a usable signal, there is no recognition of presence of unusable or foreign signals.

The *INVITATION-TO-REGISTER* messages contain identity numbers for the transmitting Access-point (API). Registered Stations know that they are hearing or not hearing the Access-point at which they are currently registered. Assuming that the Station knows or deduces that it is unregistered, the Station responds With a *REGISTER* (102) message.

Upon receipt of a *REGISTER* message, the Hub Controller decides which Access-point received that message with the best signal, and makes an assignment as described below.

This process serves Stations coming into the system from outside or from turn-ON. It does not matter if the Station moves to the coverage of

another Access-point before attempting the first message.

The infrastructure provides the *INVITATION-TO-REGISTER* (007) message at the same frequency as a complete poll described below. The message format differs from a *POLL* (007) message only in having an open destination address field. This message is sent sequentially from a group of overlapping coverage Access-points all operating on the common channel.

The unregistered Station upon hearing an *INVITATION-TO-REGISTER* sends a *REGISTER* (102) message with a long address (6 octets for LAN or possibly 60 bits for telephony within an 8 octet field) and hears immediately an *ACK* (011) and then a *PACKET-DATA-FRAME* (003) with long address in response. The payload of that frame contains the assignment of a temporary short address (2 octets) to that Station. If there is no response, the Station tries again at the next opportunity on a different Access-point. The registration response includes the identification of the Access-point from which it came, and the Station notes this as the current serving Access-point.

Normally, the response is immediate on the currently used Access-point. Sometimes the response will not be sent, and the requesting Station must wait until a new *INVITATION-TO-REGISTER* message is received before repeating the *REGISTER* message. If the *INVITATION-TO-REGISTER* message from the first used Access-point is not heard within a specified interval (e.g. 100 milliseconds) and if other Access-points are heard, a Station would reattempt registration using any Access-point with matching system identification (SYS field).

Once a Station is registered, it is periodically polled and the infrastructure knows how to reach it. The *POLL* message is the means used to be sure that each Station is present, active, assigned a short address and associated with the correct Access-point in the directory maintained by the System in the Hub Controller.

Polling Function

The infrastructure sends a *POLL* message to every known user of the system periodically. The time used is after each round of *INVITATION-TO-REQUEST* (005) messages at each of the Access-points in one reuse group when there is no pending traffic. A *POLL* round requires many rounds of *invitation* messages and could be suspended for many seconds. This time makes only a slight increase in the minimum scan time, unless there is synchronized inter-system scanning where there is no difference.

The *POLL* transmission originates on the Access-point last used by that Station, otherwise a group of surrounding Access-points is used for a second try.

The addressed Station responds with an *ACK* (110) message, and it notes the identification of the Access-point from which the *POLL* was received as current. The Station may also respond with a *REQUEST* as described later.

The air-time required for one poll is $((8+7)\text{octs}+4\mu\text{sec})$ 16 μseconds at 10 Mb/s. With 12 Stations per Access-point and 16 Access-points per group (192 Stations), the polling function uses 3.07 milliseconds per round. The frequency with which a *POLL* round is initiated is a configurable parameter.

Only the Access-points have the capacity to measure and use received signal level. If the response of a Station is at a higher signal level on a different Access-point than on the currently identified Access-point, the status entry for that Station will be changed accordingly in a system status directory after the current transaction is completed.

The Station is informed of the change by the origin API of the ACK message after sending a packet forward or by the origin API of the next transmission of a packet to the Station.

De-register Message

In response to an *INVITATION-TO-REGISTER* OR A *POLL* message, a Station should send a *DE-REGISTER* (104) message upon shut-down or

leaving the system. The polling function will eventually find this out if the message is not sent.

The opportunity for this function occurs with the same frequency as the poll described above.

Summary of Access Method for Station Originated Packets

The concept of the system is that Stations may request permission with a *REQUEST* (106/108) message to transfer data on the common channel only immediately following receiving an *INVITATION-TO-REQUEST* (005) message originated at the Hub Controller. The normal response is a *GRANT* (015) message after which the Station sends the *PACKET DATA FRAME* (114). The Hub Controller sends *ACK* (011) when and if the transmission is received without error. Alternatively, the Hub Controller can ask for a repeat using *NACK* (013) or not respond.

Peer-to-peer case: For the addressed Station to receive the message directly without repetition by the infrastructure, it must hear the *REQUEST* message, the *GRANT* message confirming and the *PACKET DATA FRAME* directly; and the Station must transmit immediate *ACK* after receiving the message before the same action by the Hub Controller.

Access-point-Originated *INVITATION-TO-REQUEST* Messages

The Hub Controller will send *INVITATION-TO-REQUEST* (005) messages only if the conditions necessary for immediate and successful transmission are present. The Hub Controller is responsible for knowing the interference possibilities that go with the use of each Access-point.

The Hub Controller originates the *INVITATION-TO-REQUEST* message consecutively at each Access-point in a reuse group to avoid interference from radio coverage overlap that might occur with simultaneous transmission.

This is one of the near indispensable functions of the infrastructure which prevents Station transmissions from interfering with transfers in progress.

Initiation of a Station-Originated Message

All Stations monitor the channel continuously and are able to hear *INVITATION-TO-REQUEST* messages usually from more than one Access-point. When a registered Station wants to send a packet, the Station listens for an *INVITATION-TO-REQUEST* message from the Access-point at which the Station was last polled or last used. At the end of the *INVITATION* message, the Station immediately sends a *REQUEST* message containing the DA, SA, SID and LEN packet header fields. There are two forms of *REQUEST* messages corresponding to long and short addressing as indicated in the TYP field.

This is a contention process with the possibility of collision made small by the frequency of opportunities to *REQUEST* and the division of the traffic so that there are only a small number of active users per Access-point.

Stations may also *REQUEST* as a second type of response to an addressed *POLL* message.

If there is no contention on the *REQUEST*, or if there is contention but a *REQUEST* message was still received successfully, the Hub Controller will send a *GRANT* message via the Access-point on which the Station responded to the *INVITATION-TO-REQUEST*. The methods for resolving contention at this step are described later below.

The Station will receive a response, *GRANT* (015) or *NACK* (013) from the addressed Access-point immediately or not at all. It is possible, and may be desirable, to have a wait interval during which the Station may receive a *GRANT*; and in this case, the infrastructure would reply with immediate *ACK* (011) (rather than *GRANT*) which would enable the wait function. The wait state would be appropriate for delays not greater than a maximum transmission length (e.g. 250 μsec).

Alternatively, there could be either *NACK* (013) which would cause *REQUEST* to be repeated or no reply to the Station *REQUEST* which would return the procedure to the

beginning. It is possible for a *REQUEST* to be rejected because the infrastructure does not have the resources to process the message at the time of the *REQUEST*.

The *GRANT* message transmission can send a power level setting, and will send the requesting Station's short address to identify the grantee. After the *GRANT* is received the Station sends the *PACKET DATA FRAME* (114).

The *GRANT* message contains a CRC-8 field on the content of the *REQUEST* message so the Station knows it is received correctly, and does not need to resend this information in the packet header. If this check fails, a new *REQUEST* is made.

After the *PACKET DATA FRAME* is sent, the originating Station waits for *ACK* (110 or 011) which may come immediately and directly from the addressed Station or slightly delayed from the Access-point. Either *ACK* ends the cycle. This is the mechanism which allows successful Station-to-Station transmission to supersede the repeat function of the infrastructure.

For virtual circuits, there is no *ACK* function since a delayed packet is a lost or useless packet. If no *ACK* is received, the Station may repeat the cycle unless the message is part of a connection-type service.

Peer-to-Peer Direct Message and ACK

The system plan permits but does not assume direct Station-to-Station communication. When an addressed Station correctly receives a message from another Station, the *ACK* (110) message is transmitted immediately. The Access-point also receives the same Station *ACK*, and then can discard its copy of the message and its intention to send a delayed *ACK* (011). If the Hub Controller does not hear the immediate *ACK* (110), it transmits a delayed *ACK* (011), and further processes the received message.

The amount of delay is not more than *ACK* message plus a propagation time, or 8 octets plus 4 μseconds (= 10.4 μsec).

Since the Hub Controller will know whether there is a possibility of direct peer-to-peer communication, it is probable that better implementations will provide the Access Manager *ACK* delay selectively. An approximate logic would be to provide the delay only for packets between stations registered on the same or immediately contiguous Access-points. If the Hub Controller knows that there is no possibility of the addressee having heard the *REQUEST* directly, the *ACK* is immediate.

For segmented peer-to-peer LAN messages, there is no possibility of auto-grant; therefore the originating Station must *REQUEST* independently for each segment after an infrastructure *INVITATION-TO-REQUEST*.

There is no provision for virtual connections directly between Stations.

Receiving Station-to-Station Direct Communication

Because of this function, it is necessary for a station to process the Station-send format message *REQUEST* (106 or 108). The *GRANT* (015) message cannot be used as an alternative because it has no provision for full source and destination addressing. There is no way for the originating Station to know that the message is being directly received from the handshake process.

If a Station hears a *REQUEST* message for which it is the addressee, it is primed for a direct transfer. The addressed Station now expects to receive the Station-send format *PACKET DATA FRAME* (114) immediately or the same frame will be retransmitted from the same or another Access-point following. This is the only circumstance where a Station processes messages from another Station rather than from an Access-point.

The direct Station-to-Station capability is configurable within the infrastructure to be active or disabled.

Resolution of Contention on *REQUESTS*

This problem is different depending on whether the contention possibilities are limited to 8 stations (as in the multi-drop ISDN S interface)

or a hundred or more that might be present in a long-reach-low-rate system.

For the case where contention is improbable on the first try, a contention type access is preferred. This can be the case when the duty cycle of any one Access-point is low--and it is unlikely to be as high as 1%.

When the contending *REQUESTS* are from Stations served on the same Access-point, it is still possible that they will be resolved without further added function.

Since the stronger radio signal masks the weaker, access may be granted in an order based on proximity rather than order-of-arrival into queue. The "work-off" rate for a few Stations and short packets is so rapid, that lack of "fairness" makes little practical difference.

It is also possible that both *REQUEST* messages may be received as contending on their home Access-point but one or the other may dominate on a second and third Access-point. It is quite possible, in the radio system, for alternate paths to resolve contention in a way having no equivalent in cable.

Nonetheless, there should be a back-up form of contention resolution to define worst-case access delay, and to obtain satisfactory operation in more defined mediums without coverage overlap.

The polling type contention recovery mode is suitable for systems with a small number of stations per access point, and is described.

The sorting method of contention recovery is suitable for perhaps 250 stations per access point, and is recognized as possible but not required.

Polling Type Contention Resolution

When the Hub Controller senses unresolvable contention at a particular Access-point, a poll of Stations registered at that Access-point is initiated. Since the number of registrants might average eight, and is rarely over 24, this consumes little time. The function is used that Stations with pending messages are allowed to *REQUEST* in response to an addressed *POLL*.

When a correct *REQUEST* message is received, the Hub Controller responds with a *GRANT* message resuming the same procedure as with handshake following the normal *INVITATION-TO-REQUEST* initiation.

Sorting Type Contention Resolution

This algorithm may be considered where there is a possibility of contention from a large number of stations sharing a common Access-point. The principle of operation is a binary sort starting with the LSB of the short address in steps requiring a message exchange using "don't care" characters.

This method is not foreseen as necessary, and therefore a detail description is not given.

Hub Controller-Originated Message

The Hub Controller can use the medium at any time by halting the sending of *INVITATION* messages. If there is waiting priority traffic via any controlled Access-point in a group, it is sent after the current *INVITATION, REQUEST, GRANT, PACKET SEND AND ACK* cycle is completed. Waiting non-priority traffic is sent after the inward cycle for that Access-point. This precedence is unimportant except when the system is heavily loaded.

From the registration procedure (and reinforced by the polling procedure), the Hub Controller knows the availability and Access-point status of each Station. Except for management controlled exceptions, no attempt would be made to reach unavailable Stations.

The Hub Controller sends a different (from the Station originated) format *PACKET DATA FRAME* (003) to the Station which has the complete header including long DA, SA, SID and LEN fields. The Station must be ready to receive these packets at any time after registering until de-registering.

At the end of the *PACKET or SEGMENT DATA FRAME* transmission, the Station sends *ACK* (110) or *NACK* (112) or nothing. There is no *ACK* for packets used for virtual circuits or for broadcast

messages. Without the *ACK* message, the Hub Controller may resend up to three tries.

The message transmitted could be from outside the network or from any Station within the network.

Signal Level Function at the Station Receiver

With this access protocol, a Station receiver is not required to measure or evaluate comparative signal level. The Station uses the Access-point assigned at registration or last used to receive a message from the infrastructure, usually a *POLL*.

SEGMENTATION AND AUTO-GRANT

There must be a limit to the maximum length of one message or connection bundle. By using a low limit for one data transfer, it is possible to allot a fraction of the capacity to each of several users rather than queue all behind a long message transmission. This is a feature that may or may not be used for packet data, but it is essential to the guarantee that a defined portion of the transmission capacity is available for connection-type services.

With limited message length, it is necessary to segment the transmission of long packets and connections. The implication is that the setup procedure is done once with full exchange of information, but thereafter segments are transmitted with only sufficient information attached for identification of the associated packet or connection.

In either case an automatically generated grant (auto-grant) procedure is used where the Hub Controller automatically sends a *GRANT* without a *REQUEST* on the channel.

The procedure for handling LAN packets longer than the protocol payload limit (e.g. 288 octets) is to divide the message into transmission segments of maximum length except for a shorter last segment. A similar procedure is used for virtual connections where each bundle of samples is processed as a segment of a message of undefined (or long) length. The marking of the last segment may be different for virtual circuits.

REQUEST-GRANT Procedure for First Segment and Setup

There is no difference in the setup procedure for a complete message or the first segment of a long message. The format is identical for the *PACKET DATA FRAME* (003) and *REQUEST* (003/108) that is a complete message or the first segment of a long message. The difference is in the content of the LEN, SID and CNN fields. These fields are not available in the *REQUEST-- SHORT ADDRESS* (106) format so this format cannot be used to initiate segmented messages.

In the 8-octet long address field, there are 4-bits set aside for distinction between LAN and ISDN addressing, and for marking first, intermediate and last segments. The definitions used are determined by future public network practice for B-ISDN, SMDS and IEEE 802.6. Similar functions are independently defined for this access protocol by the SID field in the first transmission only.

The SID field is defined at 3-bits and is always associated with LEN field of 13-bits. This size defines lengths up to 8,191 octets which is larger than the length limits in most Standard LAN protocols.

Auto-Grant

A rule of the access protocol is that a Station may not transmit except after receiving a permission message from the Hub Controller. The Hub will know that a Station has requested service for a multi-segment LAN packet from the SID and LEN fields. From these, the number of segments required can be deduced.

The first *GRANT* (015) is for the first segment only, however, with the **auto-grant** feature implemented no new *REQUEST* need be made for the following segments. The Hub Controller, using a time interval preceding that of an *INVITATION-TO-REQUEST*, can issue a *GRANT* automatically for each following segment until the transfer is complete. This avoids a need to transfer a virtual-circuit segment at an exact instant.

The addressing of following segments is short address (SA-2) as already established for Stations upon registration. The same identification is used for Station-originated segments where the short address is a pointer to the long address passed in the first *REQUEST* message. This function resembles the "Virtual Circuit Identifier" in 802.6 defined ATM cells.

Auto-grant is not required for Access-point originated messages. For transmission of segments to Stations, the Hub Controller knows when the appropriate Access-point and Station are available.

SEGMENT DATA FRAMES

The *SEGMENT DATA FRAME* uses a short address only (001/100) after the above described initiating sequences have taken place.

The SGN (segment number) counter field of 8-bits is a continuing up-counter on the number of segments transmitted which is set to 0 the first time transmitted (For Access-point originate, the second data frame and the first segment. For Station originate, the first data frame.). This counter provides a means for resequencing or for detecting missing segments. With 288 octets/payload, the value in SGN will not repeat for packets shorter than 9,790/9,216 octets (5-bits) or for a virtual circuit duration of 9 seconds.

The first 3-bits of the SGN field are used for a status indication of initial, intermediate or final segment in a long packet transfer.

It is necessary to have the CNN field in the segment header because a Station has the possibility of concurrent multiple connections during a packet transfer. A Station may transfer only one packet at a time.

COMPATIBILITY AND CAPACITY ALLOCATION FOR PACKET AND CONNECTION-TYPE TRAFFIC

With a common Access Manager and infrastructure present, absolute allocation of capacity for carried traffic is possible, because the Hub Controller determines who may transmit and the carriability of offered traffic. The logic of the Station does not participate in this choice in anyway, except for classifying the priority of its own originated traffic.

Every system or plan has a limit to the amount of traffic that can be carried. Many efficient systems carry less rather than more traffic when the level of demand reaches a critical point. When both voice and data are carried, there must be a means and method for dividing capacity in a pre-planned way so that one does not destroy the service for the other. This access protocol provides for the implementation of an adaptive or managed strategy for capacity division and handling of peak usage demands within the infrastructure and without requiring any concurrent changes in the user Station.

Excess Demand from Stations

The first method of suspending new demands for service is by withholding the *INVITATION-TO-REQUEST* message, but this cannot be the primary means because saturation of datagram and connection-type service capacity will rarely occur simultaneously.

INVITATION-TO-REQUEST will be issued, but after a *REQUEST*, *GRANT* may be withheld or *ACK* sent which orders the Station to wait. The response message adds slightly to channel loading during high traffic intervals, but the system cannot function unless the Hub Controller knows the types, size and location of waiting traffic when there is overload. When there is excess demand, the queued messages stack up in the originating Stations and not in buffer memory in the Hub Controller. The determination of the state of the available buffer memory is one of the criteria for sending the *GRANT* message.

Timely Transmission of Virtual Connection Packets

If a digital circuit is 64 kbits/second, it may be reproduced by a payload bundle of 48 octets every 6 milliseconds. The dimensional requirement is that each bundle is delivered before it is needed to have a continuous output flow from a buffer.

If the bundle dimensions are known a priori as having the these values, then the receive buffer might introduce a 3 millisecond delay so that after the first bundle is received, subsequent bundles may arrive at intervals of 6 ± 1.5 milliseconds. It can be assumed that the originating side of the wireless LAN will transmit a ready bundle within 1.5 ± 1.5 milliseconds.

In this example, the quantitizing delay is 6 milliseconds, and the transmission time uncertainty delay is a further 6 milliseconds.

The delay for speech coders now proposed for advanced digital personal and mobile telephone systems is typically 50 milliseconds. On the other hand, current Bellcore practice allows only 2 milliseconds quantitizing delay for wireless subscriber loop.

Use of Priority Function--SID field

At each use of an Access-point, the Hub Controller must first handle inward or outward connection-type messages ahead of datagrams because of the timely delivery requirement. The *REQUEST* messages contain an SID field which identifies the type of service required and the relative priority. Since connection-type services may have more than one bandwidth and gathering interval, this information is essential to capacity allocation.

The SID field is used to indicate connection-type service function, and then the LEN field is used to indicate the interval between accumulated samples and the length of the accumulated sample payload.

Segmentation for Virtual Connections

The segments of a virtual connection are treated as a segmented packet of indefinite length. The *REQUEST* announces that it is a connection-type service both in the long destination address and in the SID field, redundantly. The sampling dimensions are transmitted in the LEN field.

Each segment contains a marker that the current segment is either initial, intermediate or final in 3-bits of the SGN field.

The auto-grant initiation of the transmission of the next segment is the same as for LAN.

Segmentation Compatibility with B-ISDN

The developing broadband ISDN standards for the public network are described in Bellcore Special Report SR-NWT-001763, Issue 1, December 1990. There, and other places, the Asynchronous Transfer Mode (ATM) plan to transmit either voice or data in cells with 48 (4+44) octets of data and a 5 octet label is described. These cells may be passed at irregular intervals but at a constant average rate on a high speed medium.

This LAN protocol with its adaptive length packets and segments can conform to the ATM payload size so that a second quantitizing delay can be avoided at the boundary between an ATM based network and this wireless LAN. This flexibility would not be available with rigidly dimensioned time slotting in the wireless LAN.

CENTRAL MANAGEMENT FUNCTIONS

In a large scale wireless access system the possible use of one Access-point is in some way dependent on the status of the other Access-points around it. This is particularly true for the common channel system when there are many groups of access-points that are contiguous and part of a continuous plan. This may create a need to synchronize the use of Access-points in different groups or at least to inhibit certain use combinations. This capability is a management function of the infrastructure, and does not affect Station logic in anyway.

It is a system requirement that the access protocol contain source records for all of the data noted below, and that the transfer of information from the point of generation to the point of use be provided.

Central management is implemented mainly in the Hub Controller, but the tasks involved require support from the access protocol.

Management of Access-point Usage

To serve the functions described above, the selection of the current active Access-point takes place in an Access Manager within the Hub Controller.

Management of User Addressing, Status and Usage

A further management function is the transitory status and directory records for all active Stations with the following data:

1) Global LAN address (48 bits)
2) Local LAN address (16 bits)
3) Global E.164 address (60 bits)
4) Current Access-point identifier
5) Secondary and Tertiary Access-point identifiers
6) Current power setting
7) Last poll response time
8) Registration active/not active
9) Permitted address access
10) Changes log

The log record may be stored for some period (e.g. up to 4 weeks) following the last activity before clearing.

In addition, permanent records of usage for each active Station are needed to cover charging and cost distribution, and as evidence of system abnormality with the following data:

1) Global LAN address (48 bits)
2) Local LAN address (16 bits)
3) Global E.164 address (60 bits)
4) For each Access-point identifier used and by day and by hour--
   a) Number of messages by direction
   b) Number of octets of payload transferred
   c) Number of failed transfers requiring repeat
   d) Number and time of registrations and de-registrations including API used

Management of System Operations and Configuration

At the system central controller, a number of operational records should be kept and updated in real time. For each Access-point, the following records should be maintained:

1) Long and short Access-point identifier
2) Location and antenna pattern description
3) Hub Controller location
4) API's where channel reuse is blocked
5) API's and SYS no. of foreign systems received
6) Date of last service
7) Message handling data--
   a) Number of packet messages-- transmit, receive
   b) Number of packet octets of payload transferred
   c) Number of failed transfers requiring repeat
   d) Number of direct transfer messages
   e) Number of calls, messages and channel-seconds used for connection-type service

For the system as a whole, data should be collected on the volume and character of out-of-network traffic.

There are also configurable parameters which may require downloading from a "human" system manager. Some of these are:

1) partitioning rules for data and voice capacity
2) authorized user identifications
3) access restrictions selectively by user
4) security screening data

INTERFERENCE CONTROL BETWEEN CONTIGUOUS LANS

The presently described access protocol addresses service from a limited number of radio Access-points where there is not material overlapping radio coverage from other nearby systems. If the other nearby systems are area and capacity extensions of the first system, they are said to be commonly managed. If the other nearby systems are providing like-type service to neighboring enterprises, they are said to be independently managed; and it is for this situation that the SYS field is necessary.

Access-point Separation

Using directive antennas position for inward illumination from a boundary, it is often possible to limit overlap between Access-point transmit and receive coverage in contiguous systems to a very small portion of the served area. Also there is often geographic isolation from parking lots and other separating spaces. Nonetheless, there will be cases of inter-penetration of systems occurring in shopping malls, multi-tenant office buildings, adjacent office building, convention centers, public transportation terminals and many other places which must be considered.

The effectiveness of directional antennas at Access-points can be increased with cooperation between system managers.

Station Separation by Probability

Stations cannot use directive antennas to define their coverage. A mild consequence is that a Station in system A will hear transmissions from system B that can and must be ignored using the SYS field. A more serious consequence is that the Station in system A will not hear transmissions addressed to it because of interference from a few nearby Stations in system B.

Because each Station has an air time duty cycle which is probably less than 0.05% (3.6 seconds/hour or 18.0 megabits of transmitted data per hour at 10 Mb/s), there is only one chance in 2,000 that the interfering Station will be transmitting. Similarly, while Station A listens all the time, there is only 0.05% of the time when he is listening to addressed traffic (approximation). Only a few interfering Stations need be anticipated because there is a higher attenuation rate between Stations, and because only a few are close enough to cause interference to a particular Station.

Without considering that there is any radio frequency isolation or that the desired signal from the Access-point might be stronger than the interference, the chances of the interfering transmitter being ON while there is relevant traffic being received at the interfered Station, is about 1-in-1,000^2 between any pair of overlapping coverage Stations. If one Station had certain interference from 100 other Stations, the odds would still be 10,000:1. If fewer than 1-in-100 messages must be repeated because of interference, the system is hardly impaired.

The event of a transmitting Station originating interference to another receiving Station in a different system is statistically improbable and well within the capacity of automatic repetition to correct. The provision of automatic repetition of unacknowledged messages is a necessary part of minimizing loss from overlapping radio coverage between Stations of contiguous systems.

POWER CONTROL CONSIDERATIONS

This access protocol enables this system to operate with all transmitters at full power all of the time, but average interference levels into other systems will be reduced if power is dynamically adjusted to the level necessary.

To assist in this function, it is necessary for the Access-point receiver to measure and report the signal level of each received transmission. It is then possible for a power setting to be incorporated in the next message to the Station. Reduced power in the Station only occurs on command originating at the Hub Controller.

A further use of the report is on registration and subsequent transmissions from the Station to determine the preferred Access-point for passing messages to that Station and changes that may occur.

The power control requirement is quite different from that occurring in spread spectrum systems proposed for cellular and pocket telephone service because in this system only one Station at a time is served by each site. These other systems have many Stations simultaneously communicating with a common base Station each using a different spreading code, and it is required that all arrive with close to the same signal level.

The more precise use of output power control is an area for further study.

Communication Between Hub Controller and Access-point

Fast command transfer between Hub Controller and Access-point is needed to specify power level for the Access-point transmitter at full, -3, -6 and -9 dB or some other short set of possibilities. It is likely that power may be determined through the metallic path in the pairs between Access-point and Hub Controller rather than by reading the content of data messages.

A more difficult transmission problem will be to provide the received signal level to the Hub Controller in the reverse direction. This is also a high speed function because it is part of the delay between completion of the *REGISTER* message and the sending of the response.

BACKBONE INTERCONNECTION CONSIDERATIONS

This access protocol is independent of backbone implementation as viewed at the user Station air-interface. There are some configurations and topologies which are more favorable than others from a cost and complexity view-point.

If the propagation delay across the backbone is much longer than across wireless LAN, there may be a requirement for re-registration after roaming between LANs connected through the backbone.

CONFIGURABLE OPTIONS

There are a number of parameters which might be configurable. Table I, on the following page, is an initial effort to identify these parameters and list the possibilities. A subset of these are default values and available values required in conforming equipment.

Level A conformance is assumed to be minimum cost, and Level B is higher performance and function version.

The setting of these parameters can affect the Hub Controller, the Access-point and the Station or any combination as shown in the right column.

Line Rate

This parameter affects all major parts of the system. Initially, this rate is a configuration parameter set at a single value throughout one system; however the rate could eventually become adaptive at the Station and configurable at the Hub Controller.

Delay Intervals

The intermessage delay is configurable to provide for longer distances through radio and wire than can be accommodated with the default.

Not all systems will have a peer-to-peer direct requirement. The delay introduced with the length of one *ACK* message should not be compulsory.

TABLE I – CONFIGURABLE PARAMETERS WITH CONFORMANCE-REQUIRED VALUES

PARAMETER NAME & DESCRIPTION	PERMITTED & POSSIBLE VALUES	REQUIRED FOR CONFORMANCE			SCOPE OF EFFECT
		DEFAULT	LEVEL A	LEVEL B	
LINE RATE IN MBITS/SEC:	TBD	TBD	TBD	TBD	STN, A-P, AM
SERVICE SUPPORT	LAN, CIRCUIT	LAN	LAN	BOTH	STN, AM
INTER-MSG DELAY μSEC:	2-6	4	4, 6	2, 3, 4, 6	STN, AM
STN/STN ACK DELAY:	OFF, 4-10 OCT	8	OFF, 8	OFF, 8	STN, AM
MSG RETRIES N TIMES:	0-5	2	1, 2, 3	1, 2, 3	AM
POLL INTERVAL:	.5-10 SEC/POLL	TBD	TBD	TBD	AM
MAX SEG PAYLOAD LEN:	48-384 OCT	TBD	TBD	TBD	STN, AM
MAX PACKET LENGTH:	UP TO 10K OCT	TBD	TBD	TBD	STN, AM
DEFAULT ACCESS MGR	ON, OFF	OFF	OFF, ON	OFF, ON	STN

The interval between initiation of polls is a configurable parameter.

Lengths and Retries

If the Hub Controller misses a packet, it is a tradeoff on how many retries are appropriate.

The maximum size packet and segment are also configurable tradeoffs. Possibly, the Standard will settle on one or two values for some of these parameters.

AUTONOMOUS GROUPS NOT USING INFRASTRUCTURE

The defined area for this plan includes the words "common channel" which means not only one time-shared radio channel for both up and down link at infrastructure Access-points, but also that the same channel is used at all Access-points and communicating Stations.

The system access protocol is further designed to allow direct communication between Stations when no infrastructure is present. Two methods are possible for operation without infrastructure.

1) Access-point simulation by the first Station up
2) Distributed contention-based access method

Default Access Manager

The above described Access Protocol can include a default access manager function in each Station. When turned ON, the Station listens for *INVITATION-TO-REQUEST* messages, and hearing none acts as a reduced function Access-point by sending *INVITATION-TO-REQUEST* and *INVITATION-TO-REGISTER* messages periodically. A second Station nearby can then *REGISTER* and communicate with the first Station, and similarly for further added Stations.

The default access manager can include a prompt for manual entry of a group number with a small number of possibilities. This would be translated into the System Number field of those messages containing that field.

Distributed Contention-based Access Method

This method is similar to that described in a separate contribution relating to channelized systems (IEEE 802.1/91-96), modified to complete all setups and transfers on a single channel. In general, steps and messages are deleted which indicate and confirm channel switches.

The access method is simple. The logic is based upon the capacity to retry unsuccessful transmissions and the acknowledgement function is used to trigger retry. The same message set is used as for the infrastructure based access method previously described.

Sending a Message

If the ACK timer has expired or if no message of any kind has been received for 250 μseconds, a randomizing timer is started which adds one-of-64 possible delays randomly selected in steps of 4 μseconds. When this time has expired, the Station may transmit a *REQUEST--LONG ADDRESS* if it has a message to pass. If the addressed Station hears the *REQUEST*, then the originating Station hears the special station-originated *GRANT (115)* response. This message has the same format as the Access-point originated *GRANT* (015).

If *GRANT* is not heard, the originating Station may try twice again in quick succession for one message. No further tries may be made until the initial timing conditions again exist which permit transmission, and then only two more cycles for the same message.

If the *GRANT* response is heard, the originating Station sends the *PACKET DATA FRAME* (114). The receiving Station then sends *ACK* if received correctly or *NACK*. If *NACK* is received, the Station may resend the *PACKET DATA FRAME* immediately for two more tries otherwise the process goes back to the beginning.

Activation Criteria for Infrastructure

A composite system in which the infrastructure is dormant until needed is possible. Then the criteria for activation and dormant operation need to be defined. A possibility for these points is now described.

Quiescent infrastructure will become active sending *INVITATION* messages if any of the following conditions exist:
1) a Station asks for or is addressed by a connection-type service.
2) a LAN message is initiated which addresses a Station not also registered by a common Access-point.
3) a LAN message is received from outside addressed to a Station within the system.
4) the LAN traffic intensity results in an observed air-time usage exceeding 2% at any single Access-point within the system.

The infrastructure will not resume a quiescent condition until all of these activation conditions have been absent for a time interval of two minutes.

In the dormant state, the infrastructure will poll and solicit registration at much less frequent intervals than when it is active.

CONCLUSIONS

1. It is possible to use a single protocol for:

    A. Radio, optical and wire mediums, and

    B. Minimum, medium and high function Stations, and

    C. Short and medium distance, and

    D. Systems with a few to many hundreds of Stations, and

    E. LAN and connection-type services; provided that provision is made for:

    F. some configurable options, and

    G. a medium independent interface to allow more than one type of substitutable PHY.

2. An entirely asynchronous protocol can be evolved on the principle that the next step begins when the current step is completed; and that a high time utilization can be obtained in this way.

3. Adaptively available peer-to-peer communication can be provided.

4. All necessary functions can be obtained from a library of about 16 different message types-- half upward, half downward.

5. Protocol for operation with and without infrastructure can be incorporated in a Station with both modes using substantially the same message set and format.

6. Infrastructure might be inactive when its services are not needed.

7. The most important reasons for favoring the Common Channel Plan are:

    With infrastructure, unused capacity is in a common pool available to any Access-point in a reuse group.

    The plan is highly resistant to anomalous signal levels from contiguous Access-points as compared to channelized spread-spectrum systems.

    The plan avoids the overhead necessary for managing channel selection, and for placing this function in the ISO layer and management structure.

    It is more easily adapted to a dual mode operating with and without infrastructure.