IEEE 802.11
Wireless LAN Medium Access Control and Physical Layer Specifications

Issues Related to Distributed Time Bounded Service

11 July 1994

Kerry Lynn
Apple Computer, Inc.
One Infinite Loop, MS: 301-4J
Cupertino, CA 95014
(408) 974-2394

Larry Taylor
Apple Computer Europe, Inc.
Center de Recherche et Developpement, Le Toronto
54 route de Sartrouville
78230 Le Pecq
France
Tel: (33-1) 34 80 30 20
Fax: (33-1) 39 76 35 55

Abstract

Distributed Time Bounded Service (DTBS) can be characterized as a "best effort" attempt to deliver MSDUs within requested Quality of Service (QoS) bounds, including transit delay and delay variance. This is accomplished through priority ordering of the transmit queue and via a channel access priority mechanism. A number of issues related to DTBS are addressed.

Introduction

The 802.11 PAR requires the Standard to support a time bounded MSDU delivery service. A decision to keep Distributed Time Bounded Service in and to remove the centralized TBS from the 802.11 Standard will simplify this task and permit TBS to be made available to any station, whether a member of an ad-hoc or infrastructure network [1, 2]. This decision will also resolve positions for a number of open issues and require us to revisit a few closed issues. In some cases, the position taken on an issue may apply equally to CTBS, but in general the decision is assumed to be mutually exclusive.

Since the PAR also requires that we remain compatible with a number of approved standards including IEEE Std 802.2 LLC Sublayer/MAC Sublayer Interface Service Specification, ISO/IEC 10039 MAC Service definition, and IEEE Std 802.1d Support of the MAC Service, the positions are adapted from these standards to a large extent. The next section presents issues related to DTBS in an order that will hopefully build a better understanding of its salient aspects. Following the issues discussion, draft text is included that incorporates the positions taken.

Issues Addressed

12.5 What entities (other than LLC) will the standard support as MAC layer user ?

Alternative: Unspecified.

Argument: Since 802.11 is part of the 802 architecture, 802.11 is *required* to support 802.2 LLC and it is logical to assume that LLC will be the most likely client. However, the 802.11 charter extends only up to and including the MAC Service interface and it is impossible to predict 1) whether a MAC Service user other than LLC will be utilized, or 2) whether LLC will evolve to support new TB services. It is therefore desirable that the MAC Service provider 1) perform predictably and fairly for any client, and 2) that it provide a flexible interface capable of supporting a wide range of future services.

Position: Close with "Unspecified."

14.1 What does the support of the following connection types mean to the LLC ?
Connection without ACK,
Connectionless, and
Connection with ACK

Alternative: Beyond the scope of 802.11; refer to IEEE Std 802.2.

Position: Close by endorsing "Beyond the scope of 802.11."

12.6 What are the MAC services provided to the LLC ?

Alternative: Connectionless-mode MAC Service only.

Argument: ISO/IEC 10039 defines only one type of MAC Service; the connectionless-mode MAC Service. The MAC Service provides the following features to the MAC Service user [3]:

a) A means by which MSDUs of limited length are delimited and transparently transmitted from one source MSAP to a destination MSAP in a single MAC Service access, without establishing or later releasing a connection.

b) Associated with each instance of connectionless-mode transmission certain measures of QoS, which may be requested by the sending MAC Service user when the connectionless-mode transmission is initiated and may be modified by the MAC Service provider depending on the MAC technology.

Connectionless-mode service is required to support asynchronous MSDU delivery and is sufficient to provide DTBS provided the MAC Service user is able to associate QoS with each request and the MAC Service is able to support two or more levels of channel access priority.

Position: Close by endorsing "Connectionless-mode MAC Service only."

14.3 Where shall the connection oriented and connectionless services be integrated:,
the MAC, or
the LLC, or
somewhere else ?

Alternative: The LLC.

Argument: See 12.6. N.B. "connection oriented" is not synonymous with "time bounded."

Position: Close by endorsing "The LLC."

15.1A What does time bounded mean?

Alternative: A time bounded service, as supported by 802.11, is a "best effort" attempt to deliver a MSDU between two MSAPs within specified QoS bounds. For example, the value of a MSDU may be insensitive to a defined amount of transit delay and possibly delay variance. However when these values are exceeded, the value of the message *changes.*

Argument: See 15.1B.

Position: Close by endorsing the alternative.

15.1B What are the bounds?

Alternative: The minimum set of QoS parameters required to support DTBS include transit delay, delay variance, and user priority . The actual bounds are application specific. If the MAC Service user does not explicitly state QoS parameters, the MAC Service provider shall provide default values.

Argument: There is consensus that transit delay and delay variance are sufficient to specify a wide variety of TB services, although they are variously known by terms such as 'absolute delay' and 'variation in delay', respectively. Refer to the references for issue 15.1B.

All other QoS parameters being equal, user priority can be used to determine channel access priority of a MSDU.

Position: Close by endorsing the alternative.

12.7 What is the definition of the MAC/LLC interface for Time-bounded services ?

Alternative: IEEE Std 802.2 specifies three service primitives required of the MAC sublayer by the LLC sublayer to allow the local LLC sublayer entity to exchange LLC data units with peer LLC sublayer entities [4]:

MA-UNITDATA request
MA-UNITDATA indication
MA-UNITDATA-STATUS indication

By replacing "priority" with "QoS" in the service specification, a wide range of services can be accommodated including both asynchronous and time bounded. When the QoS that the MAC Service can be provide falls below that which was requested by the MAC Service user, the MAC Service may indicate this via MA-UNITDATA-STATUS so that (future) adaptive algorithms may react in a suitable manner. Refer to the draft text for a detailed service specification.

Position: Close by endorsing the alternative.

15.9 How will the standard address the MAC ability to service various traffic:
- Data,
- Voice, and
- Video ?

Alternative: Redundant; see issue 12.7.

Argument: Conceptual service classes may be expressed in terms of QoS parameters.

Position: Close as "Redundant; see issue 12.7."

15.11 What are the classes of Time-bounded service will the 802.11 standard specifies in addition to the required Asynchronous service ?

Alternative: Redundant; see issue 12.7.

Argument: Conceptual service classes may be expressed in terms of QoS parameters.

Position: Close as "Redundant; see issue 12.7."

15.3 What is the MAC service interface to Time Bounded Services (TBS)? and is it different from the Link Layer Connection (LLC) interface?

Alternative: Redundant; see issue 12.7.

Argument: Connection establishment is beyond the scope of 802.11. It is up to the MAC Service user to determine an a priori knowledge of the environment from the MAC Service provider and to decide whether a connection can be supported based on such knowledge. The MAC Service provider makes a best effort attempt to satisfy each request, but may reject requests that it cannot satisfy. Note that LLC generally refers to IEEE Std 802.2 - Logical Link Control, therefore this issue may be misstated.

Position: Close as "Redundant; see issue 12.7."

26.1C What is priority?

Alternative: Two types of priority are to be defined in the 802.11 MAC Specification: static (contention-free, based on duration of inter-frame spaces) and two or more levels of hierarchically independent channel access priority (contention-based; method TBD). Hierarchical independence means that increasing load from lower priority classes does not degrade the performance of higher priority classes.

Channel access priority is ultimately derived from requested QoS by a standardized algorithm. A station's channel access priority is defined as the channel access priority of its ready-to-transmit MSDU. The ready-to-transmit MSDU is defined as the MSDU at the head of the (QoS ordered) transmit queue.

Position: Close by endorsing the alternative.

26.1B Does the concept of priority need to be addressed as different traffic priorities?

Alternative: Redundant; see closure of issue 26.1C.

Position: Close with "Redundant; see closure of issue 26.1C."

19.6 What is the strategy for capacity control ?

Alternative: The strategy for partitioning of channel capacity is to prevent data from conceptually high priority user classes (voice or video) from excluding *all* data from conceptually low priority user classes (file transfer, etc.) by effectively promoting low priority data to a higher priority after sufficient queue delay. This is accomplished by means of a standardized algorithm which maps requested QoS parameters onto channel access priority.

Position: Close by endorsing the alternative.

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

Alternative: The algorithm for managing the partitioning of capacity between time bounded and asynchronous services consists of three steps:

Firstly, access fairness is determined at request time (requests that would invalidate prior QoS commitments are refused). If accepted, initial transmit queue postition is determined by a time-to-live (TTL) derived from requested QoS parameters. The TTL associated with a MSDU continues to decrement while the MSDU remains on the transmit queued.

Secondly, access fairness is ensured through a channel access priority level derived from TTL at dequeue time.

Thirdly, access fairness is determined at channel access time through the contention resolution algorithm at each channel access priority level.

Argument: Partitioning of channel capaciity is a natural consequence of the mapping of TTL to channel access priority.

Time to progress through the queue is a function of load (load = avg. pkt size * number of users) Under low load conditions what occurs is that requests that translate to a low initial TTL will tend to be transmitted before requests that map to a high initial TTL.

Under high load conditions, the liklihood of requests that translate to low initial TTL values invalidating prior QoS commitments increases. Consequently, some requests will be refused. Furthermore, under overload conditions, some queued requests will expire.

Under medium load, requests with large initial TTL values will be interleaved with short initial TTL values.

Position: Close by endorsing the alternative.

16.4 What does graceful degradation mean ?

Alternative: According to ISO/IEC 10039, a MAC Service provider may discard (unitdata) objects. A contender is a station with a ready-to-transmit MSDU. As the number of contenders increases, the channel access contention resolution protocol provides fair access to all contenders at a given channel access priority level. A natural consequence of this is that perceived capacity at each station reduces with increasing number of contenders. This orderly reduction of perceived capacity is termed "graceful degradation."

Position: Close by endorsing the alternative.

25.7 How to coordinate spectrum use between Extended Service Set (ESS)?

Alternative: Not Applicable.

Argument: No spectrum coordination is required for DTBS.

Position: Close with "Not Applicable."

Issues to be Considered

25.8 What are the implications and associated details of Clear Channel Assesment?

Alternative: DTBS relies on a channel access priority mechanism. Depending on the method adopted to provide channel access priority, DTBS may have implications for the CCA algorithm.

25.9 What Clear Channel Assesment do we put in the MAC foundation?

Alternative: See 25.8.

Issues to be Revisited

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

Alternative: The service provided by the MAC sublayer does not permit the reordering of frames transmitted with a given QoS.

Argument: Firstly, IEEE Std 802.2 contains a section on the services required of the MAC sublayer by the LLC sublayer and makes no mention of time ordering of MAC Service primitives. Secondly, ISO/IEC 10039 states:

"Only one type of object, the unitdata object, can be handed over to the MAC Service provider via a MSAP." and

"In general, the MAC Service provider may perform any or all of the following actions

a) discard objects;
b) change the order of objects.

The MAC Service exhibits a negligible rate of

a) object duplication;
b) reordering of objects for a given priority."

Thirdly, IEEE Std 802.1d states:

"The service provided by the MAC sublayer does not permit the reordering of frames transmitted with a given user priority." Clearly, the standards infer that inter-priority reordering of MSDUs may occur, and that intra-priority reordering of MSDUs must be minimized.

Current Status: CLOSE January 1994 with "No change in the order of MSDUs ..."

Position: **Reopen** and close with "The service provided by the MAC sublayer does not permit the reordering of frames transmitted with a given QoS."

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

Current Status: CLOSE January 1994 with "Yes."

Position: **Reopen** and close with "No."

Draft Text

3.2. MAC Sublayer Interface Service Specification [4].

This section specifies the services provided by the MAC sublayer. The services are described in an abstract way and do not imply any particular implementation or exposed interface.

3.2.1. Overview of Interactions

MA-UNITDATA request

MA-UNITDATA indication

MA-INITDATA-STATUS indication

3.2.2. Detailed Service Specification

3.2.2.1. MA-UNITDATA request

3.2.2.1.1. Function

This primitive defines the transfer of a MSDU from a local LLC sublayer entity to a single peer LLC sublayer entity, or multiple peer LLC sublayer entites in the case of group addresses.

3.2.2.1.2. Semantics of the Service Primitive

The semantics of the primitive are as follows:

```
MA-UNITDATA request (
                        source_address,
                        destination_address,
                        data,
                        quality_of_service,
                        service_class
                        )
```

3.2.2.2. MA-UNITDATA indication

3.2.2.2.1. Function

This primitive defines the transfer of a MSDU from the MAC sublayer entity to the LLC sublayer entity, or entities in the case of group addresses. In the absence of errors, the contents of the data parameter are logically complete and unchanged relative to the data parameter in the associated MA-UNITDATA request primitive.

3.2.2.2.2. Semantics of the Service Primitive

The semantics of the primitive are as follows:

```
MA-UNITDATA indication (
                        source_address,
                        destination_address,
                        data,
                        reception_status,
                        quality_of_service,
                        service_class
```

)

3.2.2.3 MA-UNITDATA-STATUS indication

3.2.2.3.1. Function

This primitive has local significance and shall provide the LLC sublayer user with status information for a previous associated MA-UNITDATA request primitive.

3.2.2.3.2. Semantics of the Service Primitive

The semantics of the primitive are as follows:

MA-UNITDATA-STATUS indication (
source_address,
destination_address,
transmission_status,
provided_quality_of_service,
provided_service_class
)

5.2.13. Distributed Time Bounded Service (DTBS)

An optional Distributed Time Bounded Service (DTBS) may be based on the connectionless-mode MAC Service provided by the DCF. Distributed Time Bounded Service (DTBS) can be characterized as a "best effort" attempt to deliver MSDUs within requested Quality of Service (QoS) bounds, including transit delay and delay variance..

DTBS requires a standardized algorithm to map requested Quality of Service (QoS) onto channel access priority. QoS parameters include transit delay, delay variance, and user priority. If the MAC Service user does not explicitly state QoS parameters, the MAC Service provider shall use default values. MAC Service requests that cannot be satisfied within the requested QoS bounds, or that would cause the MAC Service provider to are rejected by the MAC Service provider, thus avoiding overload conditions.

DTBS assumes that the MAC Service provides multiple hierarchically independent levels of channel access priority. Hierarchical independence means that increasing load from lower priority classes does not degrade the performance of higher priority classes.

5.2.13.1. Quality of Service

Associated with each MAC connectionless-mode transmission, certain measures of QoS are requested by the sending MAC Service user when the primitive action is initiated. The requested measures (or parameter values and options) are based on a priori knowledge by the MAC Service user of the service(s) made available to it by the MAC Service provider. Knowledge of the characteristics and type of service provided (i.e., the parameters, formats, and options that affect the transfer of data) is made available to the MAC Service user through some layer management interaction prior to (any) invocation of the MAC connectionless-mode service. Thus the MAC Service user not only has knowledge of the characteristics of the parties with which it can communicate, it also has knowledge of the statistical characteristics of the service it can expect to be provided with for each MAC Service request.

5.2.13.1.1. Transit Delay

Transit delay is the elapsed time between MA-UNITDATA.request primitives and the corresponding MA-UNITDATA.indication primitives. Elapsed time values are calculatetd only on MSDUs that are transferred successfully.

Successful transfer of a MSDU is defined to occur when the MSDU is transferred from the sending MAC Service user to the intended receiving MAC Service user without error.

For connectionless-mode transfer, transit delay is specified independently for each MAC connectionless-mode transmission. In general, satisfaction of the transit delay bound is managed by the sender.

5.2.13.1.2. Delay Variance

Delay variance is the jitter associated with transit delay. In general, satisfaction of the delay variance bound is managed by the receiver and may be used to regenerate the regular periodic interval of related sequences of MSDUs.

5.2.13.1.3. User Priority

The MAC Service user may transfer to the MAC Service provider a priori knowledge about the characteristics of the parties with which it can communicate via the user priority QoS parameter.

5.2.13.2. Mapping QoS onto Channel Priority

There is a standardized mapping of QoS Transit Delay and Delay Variance parameters to initial Time to Live (TTL). The initital transmit queue position is determined by TTL, possibly qualified by the QoS User Priority parameter. All MSDUs in the transmit queue count down their associated TTL while waiting to reach the head of the queue and be dequeued for transmission.

The channel access priority is determined, in a standardized way, from remaining TTL at dequeue time. At transmission time, the measured queue delay must be subtracted from the TTL to give the Residual Time to Live (RTL) i.e. the time left before the MSDU becomes out of date. RTL may be used in subsequent handling of the MSDU. If RTL should become less than or equal to zero, the MSDU should in all cases be discarded.

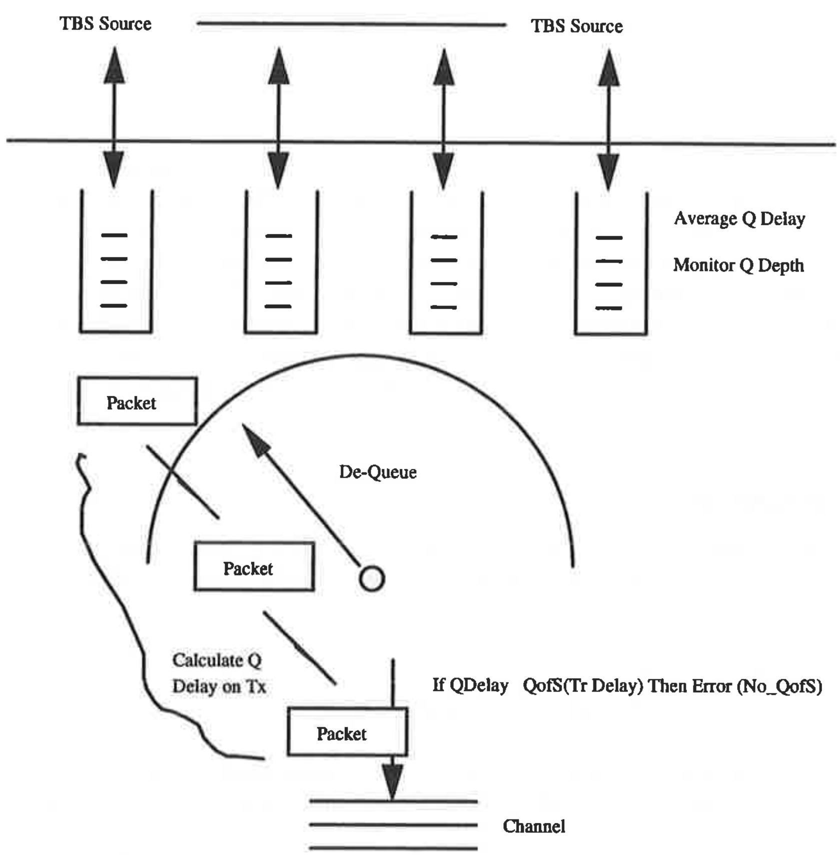

Figure ?.?: Mapping QoS onto Channel Access Priority

5.2.13.3. Partitioning of Channel Capacity

Partitioning of channel capacity amongst conceptual user classes (e.g. low priority async requests and higher priority time bounded requests) is a natural side effect of the mapping of TTL to channel access priority at dequeue time. Since *all* queued MSDUs progress towards the head of the queue as a function of their decreasing TTL, the relationship between channel access priority and conceptual user class is a function of channel load.

The strategy for partitioning of channel capacity is to prevent data from conceptually high priority user classes (voice or video) from excluding all data from conceptually low priority user classes (file transfer etc) by effectively promoting low priority data after sufficient queue delay.

The implication of this strategy to requests made with short TTL is that under high load conditions it will be known at request time that some requests may not be satisfied and will therefore be refused. Furthermore, as queue delay increases (with effective load) some requests that have been accepted may expire before being transmitted. This will favour requests which have been made with longer initial TTL (i.e. as will be likely for low priority conceptual user classes.)

The service interface may feed back to the service user the changing nature of channel access delay (via the Status primitive) as well as via the DTBS information carried with the MPDU (i.e. residual TTL and original QoS paramter values.) Applications which cannot adapt to changing channel access delay or per station capacity will suffer degradation of service. Applications which can adapt to differing conditions will be better able to continue to operate in a satsfactory manner.

In summary, capacity is partitioned in DTBS via a natural, dynamic balancing of demand and response by either refusing requests at the service interface (which would invalidate already accepted data requests) and by expiration of outdated queued requests. Capacity is therefore not partitioned statically but is allocated in a fashion which ensures that under high load conditions the channel priority mechanism is not used to permit dominatation of channel access at the expense of requests offered with less stringent QoS requirements.

5.2.13.4. Channel Access Priority Mechanism

References

[1] P. Rakity, L.Taylor, K.Lynn - Apple Computer, "A Distributed Time Bounded Service", IEEE P802.11-94/21, January 1994.

[2] T.Phipps - Symbionics, "Review of Distributed Time Bounded Services", IEEE P802.11-94/121, May 1994.

[3] ISO/IEC 10039:1991(E), Information technology - Open Systems Interconnection - Local area networks - Medium Access Control (MAC) service definition.

[4] ISO 8802-2:1989(E), IEEE Std 802.2-1989, Information processing systems - Local area networks - Part 2: Logical link control.

[5] IEEE Std 802.1d.