"Guaranteed" channel access:
Queue State element and Express traffic

Peter Johansson
Congruent Software, Inc.
PJohansson@ACM.org

Extra information not in TSPEC

- Define new “Queue state”element
  - Measures queue size in units of time, not payload
  - Simplifies scheduler in either PC or HC
- Extend categories to include “express” traffic
  - Except for external interference, ***access*** to the channel is guaranteed
  - No guarantees of channel quality

Queue state element

TID	Schedule Window	TXOP Limit	Minimum TXOP	Maximum TXOP

- TID identifies flow transmitted by a station
  - Traffic ID, “express”, ACK policy and FEC
- Schedule window (for express traffic)
- Aggregate data queued—now or periodically
- Minimum and maximum bound any single TXOP granted by the hybrid coordinator
- Traffic parameterized in time units, not bytes

Queue state element (express)

TCID	Schedule Window	TXOP Limit	Minimum TXOP	Maximum TXOP

- Schedule Window and TXOP Limit together specify the reservation
  - Arbitrarily select the start of a time period that is Schedule Window TU long
  - During that window, the station expects an aggregate grant of TXOPs equal to TXOP Limit times 16 μs
- Minimum and maximum TXOP are also measured in units of 16 μs

Queue state element (non-express)

TCID	Schedule Window	TXOP Limit	Minimum TXOP	Maximum TXOP

- Schedule Window meaningless
- TXOP Limit is the instantaneous queue size (expressed in units of 16 μs) at the station
- Minimum and maximum TXOP limit any single transmit opportunity
  - Maximum TXOP may be zero (“don’t care”)

Simplify the scheduler

- Admissions control assumed in L3 entity
- Queue sizes measured in units of 16 μs
  - Scheduler doesn't have to monitor channel quality
  - Scheduler may find it difficult to monitor channel quality in "look-aside" flows that bypass the HC
  - Scheduler doesn't have to calculate appropriate TXOP size based upon flow characteristics
- Complexity left to the end-point applications, not the scheduler!

Normative scheduler behavior

- Express traffic relies on reservations
  - L3 entity manages total time available
  - Transmitting stations are trusted to stay within the limits they have negotiated
  - In over-crowded situations, established flows continue and newcomers are denied admission
- Scheduler's obligations
  - Grant TXOPs for express traffic before other
  - Grant aggregate TXOP time reserved

"Guaranteed" channel access?

- As we all know, there is no such thing
- How may scheduler compliance with normative requirements be measured?
- Only in an ideal environment
  - Certifies correctness of the algorithm implemented
  - Does ***not*** guarantee channel access in the real world
- Normative specifications rely upon externally visible behavior, not implementation details

Contact information

Peter Johansson

Congruent Software, Inc.
98 Colorado Avenue
Berkeley, CA 94707

(510) 527-3926
(510) 527-3856 FAX

PJohansson@ACM.org

CONGRUENT SOFTWARE, INC.
98 Colorado Avenue
Berkeley, CA 94707
(510) 527-3926
(510) 527-3856 FAX

FROM: Peter Johansson

TO: IEEE P802.11 Task Group E

DATE: November 14, 2001

RE: Normative text for MAC enhancements for AV QoS

The preceding slides introduced the concept of the proposed changes; the remainder of the document contains the suggested normative text for the P802.11e working draft.

The essential components are:

a) the addition of the Queue State element to describe transmit queues at stations in terms of TIME, not bytes. The belief is that this can greatly simplify the scheduler—as well as permit the description of a category of "guaranteed channel access" (express) traffic necessary for streaming applications.

b) normative descriptions of scheduler BEHAVIOR (not implementation) with respect to express traffic.

Clock synchronization is not addressed within the draft, since it may be effected without any changes to the MAC.

7.3.2.19 Queue State element

The Queue State element contains parameters that describe the state of the transmit queue for the identified traffic category at a wireless station. The HC makes use of the information to schedule channel access during contention-free periods (including contention-free bursts). The element information field is comprised of five items as defined below and illustrated in Figure 42.15a. The total length of the information field is five octets.

Element ID (TBD)	Length (5)	TC Info (1 octet)	Schedule Window (1 octet)	TXOP Limit (1 octet)	Minimum TXOP (1 octet)	Maximum TXOP (1 octet)

Figure 42.15a – Queue State element format

The TC Info field identifies the traffic category to which the Queue State information pertains and provides additional parametric information; it is subdivided as shown in Figure 42.15b.

bits: 0	1	2	3	4	5	6	7
reserved (0)		ACK Policy	FEC	Express	Traffic category (0 – 7)		

Figure 42.15b – TC Info field

The ACK policy bit indicates whether MAC acknowledgement is required for MSDUs belonging to this traffic category and the desired form of those acknowledgements. A zero value indicates normal IEEE 802.11 acknowledgement: the addressed recipient returns an ACK or (+)CF-Ack frame after a SIFS period, according to the procedures defined in clause 9. When the ACK policy bit is one the recipient(s) shall not acknowledge the transmission.

The FEC bit is zero when transmitted data packets for this traffic category do not use forward error correction (FEC) while a value of one indicates FEC is active.

The Express bit is zero for best-effort traffic and one for express traffic. It is possible for a station to maintain separate transmit queues for the same traffic category, one for best-effort and the other for express traffic.

When the TC Info field identifies best-effort traffic, the Schedule Window field shall be zero; this indicates that the scheduler is not obligated to attempt to provide express for the traffic category. Otherwise, for express traffic, the Schedule Window field specifies the duration, in TU, of an interval used to manage channel access for the traffic category at this station. In this case, the Schedule Window and TXOP Limit together specify the amount of channel access time reserved *a priori* for the traffic category.[1]

The usage of the TXOP Limit field depends upon the value of Schedule Window. When the Schedule Window field is zero, TXOP Limit specifies the quantity of data, in units of 16 µs, buffered by the station for this traffic category. Otherwise, when Schedule Window is nonzero, the TXOP Limit field specifies the maximum aggregate TXOP time, in units of 16 µs, that should be granted the requesting ESTA during any period whose duration is equal to Schedule Window.

NOTE— Under theoretically ideal operating conditions, the TXOP Limit for express traffic is an upper limit not to be exceeded. However, when the channel has been inaccessible for some interval (perhaps due to external interference), an otherwise compliant scheduler may exceed this limit.

[1] The means by which admissions control manages the amount of theoretically available channel access time is beyond the scope of the MAC defined by this standard.

The Minimum TXOP field specifies the shortest TXOP time, in units of 16 µs, that shall be offered to the requesting ESTA when it is polled for data by the HC. The station indicates that it is unable to use a smaller TXOP.

The Maximum TXOP field specifies the longest TXOP time, in units of 16 µs, that shall be offered to the requesting ESTA when it is polled for data by the HC. The station indicates that it is unable to make use of any TXOP time offered in a single CF-Poll in excess of this limit.

Insert the following subclause, including tables and figures included therein, after 7.4.2 (which is added by P802.11e Draft 1.4), renumbering tables and figures as necessary:

7.4.3 Report Queue State frame format

The Report Queue State frame is used in a QBSS to describe the state of one or more QoS traffic queues at a wireless station that is the source of particular instances of QoS traffic. The format of the Report Queue State frame body is shown in Figure 42.16.

octets: 1	1	1	1	*n*
QoS (category 3)	**Report Queue State** (action 0)	**0**	**0**	**Queue State Element(s)**

Figure 42.17a – Report Queue State frame body

This frame is used by a station to provide the HC with current state information for one or more of the stations' Q0S transmit queues. Bandwidth management and admissions control for QoS traffic are performed by an entity whose scope is outside of the MAC; as a consequence, no response frame for the Report Queue State element is necessary.

Insert the following subclause, including tables and figures included therein, after 9.10.1.3 (which is added by P802.11e Draft 1.4), renumbering tables and figures as necessary:

9.10.1a Schedule management by the HC

When the HC performs point coordination functions it is responsible to schedule channel access for traffic categories identified in Queue State elements provide by stations within the QBSS. Individual traffic categories are either a) express or b) best-effort.[2] These traffic categories are subject to admissions control, the details of which are beyond the scope of this standard.

NOTE— Express data attempts to provide "guaranteed channel access", which phrase describes an achievable goal when the wireless medium operates free of external interference. The nature of wireless communications, particularly over unregulated frequencies, precludes absolute guarantees of channel access—however, the normative behaviors of a compliant scheduler may be observed and verified in an ideal environment that eliminates interference. Correct operation of the scheduler under ideal conditions should not be interpreted to mean that the scheduler is able to guarantee channel access under all conditions.

With express traffic (for which the scheduler attempts to provide guaranteed channel access), the normative responsibilities of the scheduler are summarized below:

a) If any express traffic is eligible to be scheduled, it shall be scheduled as soon as the HC is able to obtain channel access;

b) Higher priority traffic categories eligible to be scheduled shall be granted TXOPs before lower priority traffic categories; and

c) The scheduler shall be implemented so that, under ideal operating conditions, all stations with pending express traffic are offered TXOPs whose aggregate time is not less than TXOP Limit in any arbitrarily selected Schedule Window period.

The above requirements are behavioral requirements; they permit different scheduler implementations so long as scheduler behavior observable from the wireless medium conforms. A portion of a compliant scheduler implementation is illustrated by the pseudo-code in the table below:

```
typedef struct {
   QUEUE_STATE next;                  /* Link to another Queue State entry */
   unsigned scheduleWindow;           /* Denominated in TU (1024 ìs) */
   unsigned txopLimit;                /* Aggregate TXOP per window (units of 16 ìs) */
   unsigned minTXOP;                  /* Smallest usable TXOP (units of 16 ìs) */
   unsigned maxTXOP;                  /* Largest usable TXOP (units of 16 ìs) */
   unsigned accruedTXOP;              /* Total TXOP (units of 16 ìs) available */
   unsigned accrualInterval;          /* Time, in ìs, required to earn 16 ìs TXOP */
   unsigned leftoverTime;             /* Fractional accrual interval, in ìs */
} QUEUE_STATE;

BOOLEAN initiateCFP = FALSE;          /* Global variable available to scheduler */

/* The procedure below is invoked once every 1024 ìs. Its function is to increase the
total TXOP time accrued for each express traffic category at participating stations at
a rate calculated from the most recently received Queue State element. */

updateNextTXOP(QUEUE_STATE *queueState) {
```

[2] Best effort traffic is assumed to use a distributed contention function not managed by the scheduler at the HC.

```
   adjustedTime;                        /* Time since last TXOP accrual, in ìs */

   do {
      adjustedTime = 1024 + queueState->leftoverTime; /* Time since last invocation */
      while (adjustedTime >= queueState->accrualInterval) {
         queueState->accruedTXOP++;                  /* Accumulate another 16 ìs */
         adjustedTime -= queueState->accrualInterval;
      }
      queueState->leftoverTime = adjustedTime;       /* Save fractional remainder */
      if (queueState->accruedTXOP >= queueState->minTXOP)
         initiateCFP = TRUE;
   } while ((queueState = queueState->next) != NULL);

}
```

For each express traffic category at each participating station, the scheduler maintains queue state information described as QUEUE_STATE type above. The first four fields in the structure are transcribed directly from their corresponding fields in the Queue State element (see 7.3.2.22) each time such an element is received. The first time a Queue State element is received from a particular station for a particular express traffic category, `accrualInterval` is calculated as (Schedule Window * 1024) / TXOP Limit (rounding is permitted); the calculated value represents the number of microseconds that shall elapse before the station earns a 16 µs TXOP credit for the express traffic category. At the same time, the `accruedTXOP` field is zeroed; subsequent updates occur each time `updateNextTXOP()` is invoked.

The information stored upon receipt of a Queue State information is used by the `updateNextTXOP()` procedure, which (in this example) the scheduler invokes once every 1024 µs. This procedure manages the queue state information to provide the requested quantity of TXOPs during the schedule window defined for each station and express traffic category.[3] Whenever there is at least one express traffic category at any of the stations eligible to be granted a TXOP, the `initiateCFP` variable is set TRUE to cause a contention-free period or burst to be initiated. When the HC commences a contention-free period or burst, it shall offer TXOPs to all stations with eligible express traffic, subject to the maximum duration of a contention-free period or burst. If the HC is unable to offer TXOPs to all eligible stations, it shall initiate a subsequent contention-free period as soon as possible subsequent to the ensuing distributed contention period. A express traffic category queue at a station is eligible if `accruedTXOP` is greater than or equal to `minTXOP`, in which case the HC shall poll the station and provide a TXOP whose duration is equal to the lesser of `maxTXOP` or `accruedTXOP`. The value of `accruedTXOP` shall be decreased by the size of the TXOP offered by the CF-Poll whether or not the station makes use of the entire time.

[3] For the sake of simplicity, the procedure shown is applicable to a single traffic category but is easily generalized to multiple traffic categories.

Insert the following subclause, including tables and figures included therein, after 10.3.11 (which is added by P802.11e Draft 1.4), renumbering tables and figures as necessary:

10.3.11a Queue state update

The following primitives describe how queue state is modified within a QBSS.

10.3.11a.1 MLME-QUPDATE.request

10.3.11a.1.1 Function

This primitive requests update of queue state.

10.3.11a.1.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-QUPDATE.request (WSTAAddress,
TC,
Queue State
)

Name	Type	Valid Range	Description
WSTAAddress	MACAddress	any valid MAC address	Specifies the MAC address of the WSTA that is the context of the traffic category for which the queue state is being updated.
TC	Integer	0-7	Specifies the traffic category for which the queue state is being updated. Note that the MA-UNITDATA primitives that handle MSDUs for the specified queue use a priority parameter value of (TC+8).
Queue State	As defined in Frame Format	As defined in Frame Format	The parameter values that specify the QoS for the designated traffic category.

10.3.11a.1.3 When generated

This primitive is generated by the SME at an ESTA when a higher-layer QoS management entity wishes to update queue state.

10.3.11a.1.4 Effect of receipt

This primitive initiates a queue state update procedure. The MLME subsequently issues a MLME-QUPDATE.confirm that reflects the results.

10.3.11a.2 MLME-QUPDATE.confirm

10.3.11a.2.1 Function

This primitive reports the results of a queue state update attempt.

10.3.11a.2.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-QUPDATE.confirm (ResultCode
)

Name	Type	Valid Range	Description
ResultCode	enumeration	SUCCESS, INVALID_ PARAMETERS, TIMEOUT	Indicates the result of the MLME-QUPDATE.request.

10.3.11a.2.3 When generated

This primitive is generated by the MLME as a result of an MLME-QUPDATE.request to update state information for a specified queue within the QBSS.

10.3.11a.2.4 Effect of receipt

The SME is notified of the results of the queue state element update procedure.

10.3.11a.3 MLME-QUPDATE.indication

10.3.11a.3.1 Function

This primitive reports the occurrence of an update to queue state within the QBSS at the HC.

10.3.11a.3.2 Semantics of the service primitive

The primitive parameters are as follows:

MLME-QUPDATE.indication (WSTAAddress,
TC,
Queue State
)

Name	Type	Valid Range	Description
WSTAAddress	MACAddress	any valid MAC address	Specifies the MAC address of the WSTA that is the context of the traffic category for which queue state is being updated.
TC	Integer	0-7	Specifies the traffic category for which queue state is being updated. Note that the MA-UNITDATA primitives that handle MSDUs for the specified queue use a priority parameter value of (TC+8).
Queue State	As defined in Frame Format	As defined in Frame Format	The parameter values that specify the QoS for the designated traffic category.

10.3.11a.3.3 When generated

This primitive is generated by the MLME as a result of the receipt of a Report Queue State QoS action frame.

10.3.11a.3.4 Effect of receipt

The SME is notified of the occurrence of the queue state update procedure and action and virtual stream parameters for that update.