

## 10. Transmit state machines (proposal 2)

NOTE—Multiple bunch-avoiding pacing protocols are presented for consideration:  
 a) Clause 9 presents a pseudo-synchronous transmission model.  
 b) Clause 10 (this clause) presents cross-flow shaper transmission model.

### 10.1 Rate-based scheduling overview

The clause describes a rate-based scheduling technique. The rate-based scheduling concepts are similar to those within rate monotonic scheduling protocols, commonly used within real-time systems. Objectives associated with time-sensitive forwarding alternatives include the following:

- a) Multiple time-sensitive transmission rates are supported, including:
  - 1) Highest rate 8 kHz traffic, such as the traffic generated by simple bridges between RE and existing IEEE 1394[B6] A/V devices.
  - 2) Lower rate traffic, such as voice over internet protocol (VOIP) traffic, without forcing this traffic to be reblocked into smaller (and therefore less efficient) frame sizes.
- b) Frame forwarding should not be dependent on successful time-of-day synchronization between the bridge and adjacent stations. Frame forwarding should succeed before the grand clock-master station has been selected, or when the selected grand-master clock station changes.
- c) Frame-forwarding protocols should leverage existing bridge queue and service models, although specification of abstract rate shaper details is expected.

Rate-based scheduling involves associating a priority with frame transmissions, where the priority is a monotonic function of the frame transmission frequency, as illustrated in Figure 10.1. Assuming the cumulative traffic is limited to less than the link capacity, the latency of each traffic class is guaranteed (the latency guarantee is approximately an MTU more than an inter-arrival period).




Figure 10.1—Rate-based priorities

**10.1.1 Rate-based priorities**

Quality of service is based on the availability of user\_priority field parameter associated within transmitted time-sensitive frames, as listed in Table 10.1.


**Table 10.1—Tagged priority values**

Code	Interval (ms)	Name	Description
0	n/a	CLASS_C	Best effort, with minimal guaranteed BW
1	n/a	CLASS_B	Preferred, with minimal guaranteed BW
2	—	—	Used for other purposes
3	8	CLASS_A3	Guaranteed BW over longer interval
4	2	CLASS_A2	Guaranteed BW over long interval
5	0.5	CLASS_A1	Guaranteed BW over short interval
6	0.125	CLASS_A0	Guaranteed BW over shorter interval
7	—	—	Used for other purposes


1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54

**10.1.2 Port-to-port reshaping, without bunching accumulation**

The concept of rate-based scheduling assumes shaped talkers and reshaped talker agents within bridges, as illustrated in Figure 10.2 (only the components associated with specific flows are illustrated). In this illustration, classA0 traffic flows between points (a, b, c, d, e), exhibits bunched and reshaped behaviors, as illustrated in Figure 10.3.



**Figure 10.2—Reshaped bridge-traffic topology, with bunching control**



**Figure 10.3—Reshaped bridge-traffic timing**

The (a) through (e) time lines represent the flow of frames from within one talker-or-bridge into another bridge-or-listener, described as follows:


- a) A properly shaped source stream is originally generated within a talker, or a port-to-port flow (consisting of multiple streams) within a bridge.
- b) Forwarding of multiple sources to a shared transmission link can produce jitter, due to slight differences in frame-to-frame spacings.
- c) Forwarding of multiple sources to a shared transmission link can produce additional jitter, when higher-class traffic waits for the completion of previously initiated lower-class transmissions.
- d) Bunching becomes apparent in the port-to-port flow, representing the portion of the received (c) traffic that is forwarded to a specific transmitter port.
- e) A shaper delays the forwarding of bunched frames, so that the port-to-port flow is properly shaped. Delays can be invoked by time stamping frames with an in-the-future transmission time.

The reshaped flow (e) retains the properly shaped properties of the preceding flow (a), while incurring a maximum delay  $d$  through the bridge. These properties ensure a linear maximum delay of  $n \times d$ , for streams that flow through  $N$  bridges.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54

**10.1.3 Port-to-port reshaping, with bunching control**

The complexity of managing traffic classes can be reduced by eliminating shapers, so that receiver inputs are merged before being shaped, as illustrated in Figure 10.4 (only the components associated with specific flows are illustrated). In this illustration, classA0 traffic flows between points (a, b, c, d, e), exhibits bunched and reshaped behaviors, as illustrated in Figure 10.3.



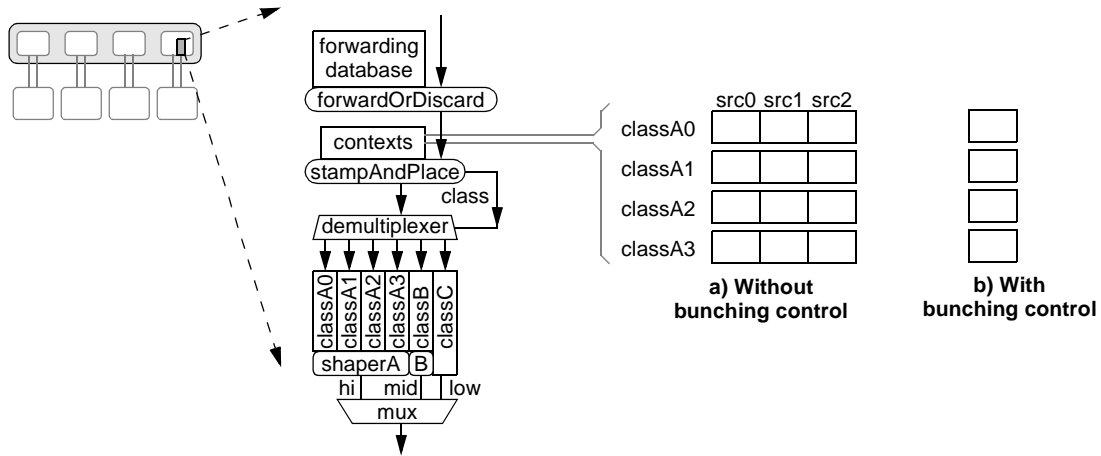
**Figure 10.4—Reshaped bridge-traffic topology, without bunching control**

This cost-reduced approach solves the bursting problem, without addressing the bunching problem. Within constrained topologies this can sometimes be sufficient to meet worst-case latency requirements.

**10.1.4 Transmit ports**

**10.1.4.1 Transmit port structure**

The transmit port is responsible for shaping classA traffic (to avoid bunching) and pacing classB/classC traffic (to avoid classC traffic starvation). Pacing and shaping algorithms assume functionally distinct queues within each transmit port, as illustrated in Figure 10.5.



**Figure 10.5—Transmit-queue structure**

The intent of per-class shapers is to avoid priority inversions, wherein higher-class frames are delayed by the presence of concurrent lower-class traffic. Independent per-class shapers and queues allow enqueued higher-class and lower-class frames to be forwarded independently, thus avoiding priority inversions within queues.

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54

The intent of per-source shapers is to avoid increasingly large cumulative bunching delays. The per-source reshaping eliminates bunches before merging, so that the pass-through bunching severity for 1-bridge and  $n$ -bridge flows are the same.

#### 10.1.4.2 Enqueue reshaping contexts

The desired per-class latencies could not be guaranteed in the presence of classA traffic bunching. To avoid bunching, frames are shaped before being placed into classified transmit queues.

A shaper is responsible for attaching a time-stamp label to frames. With bursting control, a time-stamp shaper is logically associated with each source port and each classA traffic subclass (classA0, classA1, classA2, classA3). E.g, a four-port switch (which has three possible source ports) would have 12 time-stamp shapers on each transmit port. Without bursting control, a time-stamp shaper is logically associated with each classA traffic subclass (classA0, classA1, classA2, classA3). E.g, a N-port switch (which has N-1 possible source ports) would have 4 time-stamp shapers on each transmit port.

The purpose of a time-stamp shaper is to associate a time-stamp label with each queued frame. The time-stamp label represents a time in the future; the frame's transmission is deferred until the current time reaches the frame's time-stamp value. This facilitates the delayed forwarding of successive frames within each bunch, thus suppressing the bunching effects found on receive-link transmission.

The context for each time-stamp shaper is based on the frame's receive port and traffic class. While the context is considerably larger than that associated with strict per-port shapers, only one shaper (within each port) is ever active. Thus, context-switching per-port shaper instances represent a viable implementation technology

### 10.1.4.3 Dequeue shaping and pacing

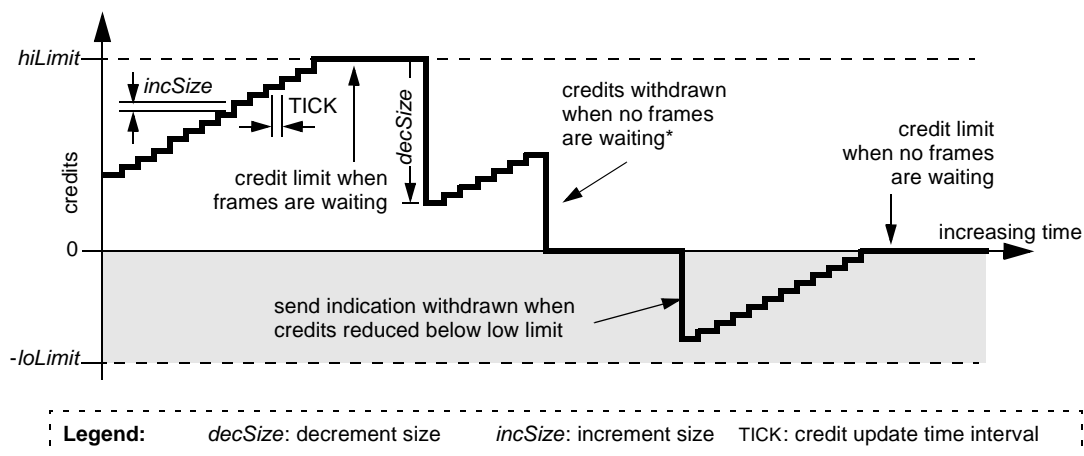
Transmit ports utilize a shaper and pacer, as illustrated as shaperA and B components within Figure 10.5. The purpose of these is to ensure forward progress of best-effort control traffic. In concept, this involves a two-step bandwidth partitioning mechanism:

- The shaperA limits the cumulative classA and primary classB traffic to 75% of the link bandwidth. The intent is to ensure that 25% residual bandwidth remains available for lower-class traffic.
- Pacer B partitions the residual 25% traffic equally between classB and classC traffic. This ensures that classB traffic is never starved, in the presence of 75% classA traffic. This ensures that classC traffic is not starved, in the presence of excess classB traffic.

### 10.1.5 Credit-based shapers and pacers

#### 10.1.5.1 Credit-based shapers

Although multiple shapers are specified within this working paper, the behavior of most shapers can be characterized by a common algorithm and instance-specific parameters (as done within RPR[B5]). The shaper's credits are adjusted down or up, as illustrated in Figure 10.6. The decrement and increment values typically represent sizes of a transmitted frame and of credit increments in each update interval, respectively.



**Figure 10.6—Credit-based shapers**

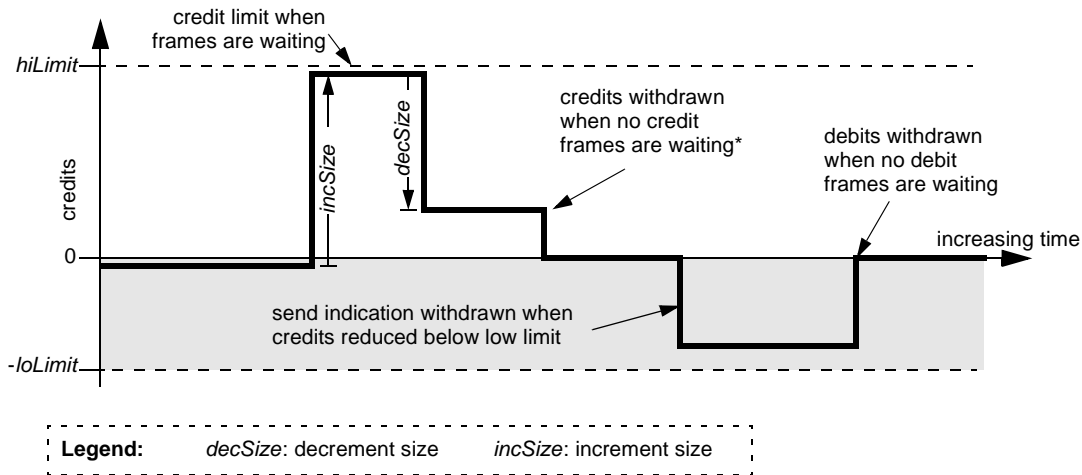
Crossing below the zero threshold generates a rate-limiting indication, so that offered traffic can stop. By design, the credit value never goes below the  $-loLimit$  extreme. To bound the burst traffic after inactivity intervals, when no frames are ready for transmission, credits are reduced to zero (if currently higher than zero) and can accumulate to no more than the zero-value limit.

The  $hiLimit$  threshold limits the positive credits, to avoid overflow. When frames are ready for transmission (and are being blocked by transit traffic), credits can accumulate to no more than this  $hiLimit$  value.

In concept, the shaper consist of a token bucket. The number of credits in a token bucket is decremented by the size of each transmitted frame. The credits in the token bucket are incremented at the end of every credit update interval. A frame is only transmitted when the credits are positive.

### 10.1.5.2 Credit-based pacers

Although multiple pacers are specified within this working paper, the behavior of most pacers can be characterized by a common algorithm and instance-specific parameters (as done within RPR[B5]). The pacer's credits are adjusted down or up, as illustrated in Figure 10.7. The decrement and increment values typically represent sizes of debit and credit frames in each update interval, respectively.



**Figure 10.7—Pacer credit adjustments over time**

Crossing below the zero threshold generates a rate-limiting indication, so that offered traffic can stop. By design, the credit value never goes below the  $-loLimit$  extreme. To bound the burst traffic after inactivity intervals, when no frames are ready for transmission, credits are reduced to zero (if currently higher than zero) and can accumulate to no more than the zero-value limit.

The  $hiLimit$  threshold limits the positive credits, to avoid overflow. When frames are ready for transmission (and are being blocked by transit traffic), credits can accumulate to no more than this  $hiLimit$  value.

In concept, the pacer consists of a token bucket. The credits in the token bucket are incremented by the size of each transmitted debit-frame. The number of credits in a token bucket is decremented by the size of each transmitted credit-frame. A credit-frame is only transmitted when the credits are positive; a debit-frame is only transmitted when the credits are negative.

## 10.2 Terminology and variables

### 10.2.1 Common state machine definitions

The following state machine inputs are used multiple times within this clause.

queue values	1
Enumerated values used to specify shared queue structures.	2
QP_TX_PUSH—The transmit port’s internal queue, where received frames are placed.	3
QP_TX_A0—The first of the output port’s classA buffers.	4
QP_TX_A1—The second of the output port’s classA buffers.	5
QP_TX_A2—The third of the output port’s classA buffers.	6
QP_TX_A3—The fourth of the output port’s classA buffers.	7
QP_TX_A4—The second of the output port’s classA buffers.	8
QP_TX_A5—The third of the output port’s classA buffers.	9
QP_TX_BP—The output port’s classB queue.	10
QP_TX_CP—The output port’s classC queue.	11
QP_TX_LINK—The output port’s transmit-PHY queue.	12

## 10.2.2 Common state machine variables

One instance of each variable specified in this clause exists in each port, unless otherwise noted.

*currentTime*

A value representing the current time.

*framed*

The contents of a received frame, with supplemental information, as follows:

*frame*—The contents of a frame.

*sourcePort*—The source port that received the frame.

*txTime*—A time-stamp value representing the intended (bunching delayed) transmission time.

## 10.2.3 Common state machine routines

*Max(value1, value2)*

Returns the numerically larger of two values.

## 10.2.4 Variables and routines defined in other clauses

This clause references the following variables and routines defined in Clause 7:

*currentTime*

See 7.2.2.

*Dequeue(queue)*

*Enqueue(queue, frame)*

*Min(value1, value2)*

See 7.2.3.

## 10.3 Pacing state machines

### 10.3.1 TransmitRx state machine

The TransmitRx state machine is responsible for enqueueing traffic (received on other ports and broadcast to all possible transmitter ports) for possible forwarding. An intent is to transfer each to the appropriate output queue.

The following subclasses describe parameters used within the context of this state machine.



<b>10.3.1.1 TransmitRx state machine definitions</b>	1
queue values	2
Enumerated values used to specify shared queue structures.	3
QP_TX_A0, QP_TX_A1, QP_TX_A2, QP_TX_A3, QP_TX_A4, QP_TX_A5	4
QP_TX_BP, QP_TX_CP	5
QP_TX_PUSH	6
See 10.2.1.	7
	8
	9
<b>10.3.1.2 TransmitRx state machine variables</b>	10
	11
<i>class</i>	12
A value that represents the frame's priority class.	13
<i>count</i>	14
A value that represents the current credits, while minimum and maximum limits are being applied.	15
<i>currentTime</i>	16
See 10.2.4.	17
<i>delay</i>	18
A value that represents the time delay assigned by the frame's shaper.	19
<i>framed</i>	20
See 10.2.2.	21
<i>sPtr</i>	22
Represents a pointer to shaper values.	23
	24
<b>10.3.1.3 TransmitRx state machine routines</b>	25
	26
<i>ContextCheck(sourcePort, class)</i>	27
Returns a pointer to the associated pacer context, with the following fields:	28
<i>credit</i> —The cumulative credit from past pacer activities.	29
<i>lastTime</i> —The last time the pacer was invoked.	30
<i>loLimit</i> —The low limit for shaper credits.	31
<i>rate</i> —The highest allowed rate of the paced traffic, in bytes-per-second.	32
<i>Dequeue(queue)</i>	33
See 10.2.4.	34
<i>Enqueue(queue, frame)</i>	35
Places the <i>frame</i> at the tail of the specified <i>queue</i> within the assumed port.	36
<i>ForwardClass(framed)</i>	37
The forwarding database is checked. If forwarding is enabled, the priority class is returned.	38
Otherwise, a NULL class value is returned. The following enumerated values are returned:	39
CLASS_A0—The associated multicast frame is forwarded as classA traffic.	40
CLASS_A1—The associated multicast frame is forwarded as classA traffic.	41
CLASS_A2—The associated multicast frame is forwarded as classA traffic.	42
CLASS_A3—The associated multicast frame is forwarded as classA traffic.	43
CLASS_A4—The associated multicast frame is forwarded as classA traffic.	44
CLASS_A5—The associated multicast frame is forwarded as classA traffic.	45
CLASS_B—The associated multicast frame is forwarded as classB traffic.	46
CLASS_C—The associated multicast frame is forwarded as classC traffic.	47
<i>Max(value1, value2)</i>	48
See 10.2.3.	49
<i>Min(value1, value2)</i>	50
See 10.2.4.	51
	52
	53
	54

**10.3.1.4 TransmitRx state table**

The TransmitRx state machine is specified in Table 9.2. In the case of any ambiguity between the text and the state machine, the state machine shall take precedence. The notation used in the state table is described in 3.4.

**Table 10.2—TransmitRx state table**

Current		Row	Next	
state	condition		action	state
START	(framed = Dequeue(QP_TX_PUSH))!=NULL	1	—	FIRST
	—	2	—	START
FIRST	(class = ForwardClass(framed)) != NULL	3	sPtr = ContextCheck(framed.sourcePort, class);	NEXT
	—	4	—	START
NEXT	class == CLASS_A0	5	queue = QP_TX_A0;	PACE
	class == CLASS_A1	6	queue = QP_TX_A1;	
	class == CLASS_A2	7	queue = QP_TX_A2;	
	class == CLASS_A3	8	queue = QP_TX_A3;	
	class == CLASS_A4	9	queue = QP_TX_A4;	
	class == CLASS_A5	10	queue = QP_TX_A5;	
	class == CLASS_B	11	queue = QP_TX_BP;	
	—	12	queue = QP_TX_CP;	
PACE	—	13	count = sPtr->credit + sPtr->rate * (currentTime - sPtr->lastTime) - Size(frame); count = Max(0, Min(sPtr->loLimit, count)); sPtr->credit = count; sPtr->lastTime = currentTime; delay = Max(0, -count / sPtr->rate); framed.txTime = currentTime + delay;	FINAL
FINAL	—	14	Enqueue(queue, frame);	START

**Row 10.2-1:** If a frame has arrived, process that frame.

**Row 10.2-2:** Otherwise, wait for the next frame to arrive.

**Row 10.2-3:** When forwarded frames, the shaper context is based on the source port and class.

**Row 10.2-4:** The non-forwarded frames are discarded.

**Row 10.2-5:** The classA0 frames are forwarded to the appropriate time-sensitive classA0 queue.

**Row 10.2-6:** The classA1 frames are forwarded to the appropriate time-sensitive classA1 queue.

**Row 10.2-7:** The classA2 frames are forwarded to the appropriate time-sensitive classA2 queue.

**Row 10.2-8:** The classA3 frames are forwarded to the appropriate time-sensitive classA3 queue.

**Row 10.2-9:** The classA4 frames are forwarded to the appropriate time-sensitive classA4 queue.

**Row 10.2-10:** The classA5 frames are forwarded to the appropriate time-sensitive classA5 queue.

<b>Row 10.2-11:</b> The classB frames are forwarded to the appropriate time-sensitive classB queue.	1
<b>Row 10.2-12:</b> The classC frames are forwarded to the appropriate time-sensitive classC queue.	2
	3
<b>Row 10.2-13:</b> ClassA frames are time stamped by shapers.	4
Shaper pacer parameters for each distinct { <i>class</i> , <i>source</i> } pair constrains bunching within each class.	5
Shaper parameters are updated as in 10.1.5.1: decremented on transmissions and incremented over time.	6
High and low limits are applied to the updated credits and the last-updated time is updated.	7
Negative credits correspond to transmission-delay values, which are attached to output-port queued frames.	8
	9
<b>Row 10.2-12:</b> The received frames are placed into the appropriate queue.	10
	11
<b>10.3.2 TransmitTx state machine</b>	12
	13
The TransmitTx state machine is responsible for pacing/shaping classA traffic and shaping classB traffic destined for 1 Gb/s links. An intent is to support projected MTU-sized transfers and interleaved lower-class traffic, without exceeding the 1-cycle delay inherent with cycle-synchronous bridge-forwarding protocols.	14
	15
	16
The following subclasses describe parameters used within the context of this state machine.	17
	18
	19
<b>10.3.2.1 TransmitTx state machine definitions</b>	20
	21
BPS	22
The nominal link transmission rate, in bytes per second.	23
MTU	24
The maximum frame size, in bytes.	25
queue values	26
Enumerated values used to specify shared queue structures.	27
QP_TX_A0, QP_TX_A1, QP_TX_A2, QP_TX_A3	28
QP_TX_BP, QP_TX_CP	29
QP_TX_LINK	30
See 10.2.1.	31
TICK	32
The amount of time between shaper updates.	33
Range: [1 bytes transmit time, 16-bit transmit time]	34
Default: 1 byte transmit time	35
	36
<b>10.3.2.2 TransmitTx state machine variables</b>	37
	38
<i>best</i>	39
A value that represents the weight and identify of the next-best classA queue.	40
<i>goodness</i> —The smallest <i>weight</i> × <i>wait</i> value associated with alternate classA transmissions.	41
<i>queue</i> —The queue associated with the best futuristic encapsulated frame.	42
<i>countA</i>	43
A speculative value of creditA, used only when the frame is qualified for transmission.	44
<i>countB</i>	45
A speculative value of creditB, used only when the frame is qualified for transmission.	46
<i>creditA</i>	47
A shaper credit whose positive value enables classA/classB primary transmissions.	48
<i>creditB</i>	49
A shaper credit value whose positive and negative values enable secondary classB and classC transmissions respectively.	50
	51
<i>currentTime</i>	52
See 10.2.4.	53
	54

<i>frame</i>	1
The contents of a to-be-transmitted frame.	2
<i>framed</i>	3
See 10.2.2.	4
<i>hiLimitA</i>	5
A value that limits the cumulative <i>creditA</i> credits.	6
Value: MTU.	7
<i>hiLimitB</i>	8
A value that limits the cumulative <i>creditB</i> credits.	9
Value: MTU.	10
<i>limit</i>	11
A value that limits the amount of transmitted primary classA/classB bandwidth.	12
<i>loLimitA</i>	13
A value that limits the cumulative <i>creditA</i> debits.	14
Value: MTU.	15
<i>loLimitB</i>	16
A value that limits the cumulative <i>creditB</i> debits.	17
Value: MTU.	18
<i>tickTime</i>	19
A value that defines when the time-tick interval ends.	20
<b>10.3.2.3 TransmitTx state machine routines</b>	21
	22
<i>Dequeue(queue)</i>	23
See 10.2.4.	24
<i>Unqueue(queue, weight, &amp;best, currentTime)</i>	25
Dequeues and returns the most overdue frame from the specified <i>queue</i> , excluding those frames	26
whose scheduled transmission time is after the specified <i>currentTime</i> value.	27
<i>framed</i> —The oldest of the overdue frame.	28
NULL—No frame available.	29
In the presence of only futuristic frames, a <i>test</i> = <i>weight</i> ×( <i>txTime</i> - <i>currentTime</i> ) value is computed.	30
If <i>best.queue</i> is NULL or <i>test</i> < <i>best.goodness</i> , the <i>best.queue</i> and <i>best.goodness</i> components are	31
updated to reflect the best alternate classA transmission queue.	32
<i>Enqueue(queue, frame)</i>	33
See 10.2.4.	34
<i>Size(frame)</i>	35
Returns the size of the specified frame.	36
<i>StaleFrame(frame, queue)</i>	37
Indicates whether the specified frame is stale and discardable, as specified by Equation 10.1.	38
0—The specified frame is not stale.	39
1—(Otherwise.)	40
	41
// The value of "internal" depends on the class, as specified in Table 10.1. (10.1)	42
(index = (QP_TX_A0 - queue),	43
(currentTime - framed.txTime) > (2 * (MTU + interval[index] * BPS)))	44
	45
	46
	47
	48
	49
	50
	51
	52
	53
	54

**10.3.2.4 TransmitTx state table**

The TransmitTx state machine is specified in Table 9.3. In the case of any ambiguity between the text and the state machine, the state machine shall take precedence. The notation used in the state table is described in 3.4.

**Table 10.3—TransmitTx state table**

Current		Row	Next	
state	condition		action	state
START	$(currentTime - tickTime) \geq TICK;$	1	$creditA = \text{Min}(hiLimitA, creditA + 0.75 * TICK * BPS);$ $tickTime = currentTime;$	START
	$!QueueEmpty(QP\_TX\_LINK)$	2	—	
	$creditA < 0$	3	—	FAIR
	—	4	$best.queue = NULL;$	BEST
BEST	$(framed = Unqueue(queue= QP\_TX\_A0, 32, \&best, currentTime)) \neq NULL$	5	$countA = \text{Min}(loLimitA, creditA - \text{Size}(framed));$	NEAR
	$(framed = Unqueue(queue= QP\_TX\_A1, 16, \&best, currentTime)) \neq NULL$	6		
	$(framed = Unqueue(queue= QP\_TX\_A2, 8, \&best, currentTime)) \neq NULL$	7		
	$(framed = Unqueue(queue= QP\_TX\_A3, 4, \&best, currentTime)) \neq NULL$	8		
	$(framed = Unqueue(queue= QP\_TX\_A4, 2, \&best, currentTime)) \neq NULL$	9		
	$(framed = Unqueue(queue= QP\_TX\_A5, 1, \&best, currentTime)) \neq NULL$	10		
	$best.queue \neq NULL \ \&\& \ (framed = Dequeue(queue= best.queue)) \neq NULL$	11		
	$(framed = Dequeue(QP\_TX\_BP)) \neq NULL$	12		
	—	13	$creditA = 0;$	START
FAIR	$creditB \geq 0 \ \&\& \ (framed = Dequeue(QP\_TX\_BP)) \neq NULL$	14	$creditB = creditB - \text{Size}(framed);$	FINAL
	$creditB \leq 0 \ \&\& \ (framed = Dequeue(QP\_TX\_CP)) \neq NULL$	15	$creditB = creditB + \text{Size}(framed);$	
	$(framed = Dequeue(QP\_TX\_BP)) \neq NULL$	16	$creditB = 0;$	
	$(framed = Dequeue(QP\_TX\_CP)) \neq NULL$	17		
	—	18	$creditB = 0;$	START
NEAR	$\text{StaleFrame}(framed, queue)$	19	—	START
	—	20	$creditA = countA;$	FINAL

1  
2  
3  
4  
5  
6  
7  
8  
9  
10  
11  
12  
13  
14  
15  
16  
17  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45  
46  
47  
48  
49  
50  
51  
52  
53  
54

**Table 10.3—TransmitTx state table (continued)**

Current		Row	Next	
state	condition		action	state
FINAL	—	21	Enqueue(QP_TX_LINK, framed.frame);	START

**Row 10.3-1:** Update the classA credits after each tick interval.

**Row 10.3-2:** Wait for the queue to be emptied, so that something can be transmitted.

**Row 10.3-3:** In the absence of classA credits, fairly transmit enqueued classB and classC frames.

**Row 10.3-4:** Fairly service classB/classC when the classA/classB transmissions are disallowed.

**Row 10.3-5:** If enabled and available, a classA0 frame is transmitted.

**Row 10.3-6:** If enabled and available, a classA1 frame is transmitted.

**Row 10.3-7:** If enabled and available, a classA2 frame is transmitted.

**Row 10.3-8:** If enabled and available, a classA3 frame is transmitted.

**Row 10.3-9:** If enabled and available, a classA4 frame is transmitted.

**Row 10.3-10:** If enabled and available, a classA5 frame is transmitted.

**Row 10.3-11:** If available, a scheduled-for-the-future classA frame is transmitted.

**Row 10.3-12:** If enabled and available, a classB frame is transmitted.

**Row 10.3-13:** Since nothing is ready to be sent, the classA credits are cleared.

**Row 10.3-14:** If enabled and available, a classB frame is transmitted.

The *creditB* values is decremented by the transmitted frame size, to avoid classC starvation.

**Row 10.3-15:** If enabled and available, a classC frame is transmitted.

The *creditB* values is incremented by the transmitted frame size, to avoid classB starvation.

**Row 10.3-16:** If available, a classB frame is transmitted.

**Row 10.3-17:** If available, a classC frame is transmitted.

**Row 10.3-18:** Otherwise, no frame is transmitted.

**Row 10.3-19:** Stale frames, whose delivery times cannot be guaranteed, are discarded.

**Row 10.3-20:** Non-stale frames are not discarded.

**Row 10.3-21:** The next frame is transmitted and credits are updated accordingly.