

9. Fairness

9.1 Fairness scenarios

This subclause attempts to show that the current fairness protocols have not adequately addressed the possible effects of classC traffic on (intended to be) higher precedence classA and classB traffic. The subannex is formatted for inclusion in Clause 9, for the convenience of the reader. This contribution editor is primarily concerned with clarity, correctness, and availability. The final placement of this content could be Clause 9, an annex, or a working-group presentation.

This subclause is intended to illustrate fairness concerns, particularly the compromise of classA and classB guarantees resulting from the dependence on fairness-eligible protocols dynamics. Discussion of the following problems will be provided:

- a) ClassA0. The classA0 concern is related to the downstream shaper, discussed in 9.2.
 - 1) The problem of classA0 blockage by upstream STQ traffic is discussed in 9.2.1. The problem is that cumulative STQ depths contribute to the worst-case jitter.
 - 2) The solution of applying a downstream shaper to non-classA traffic is discussed in 9.2.3. The solution is applies the downstream shaper to add/transit classB/C traffic.
- b) ClassA1. The problem with classA1 traffic is related to the management of STQ buffers (TBD).
- c) ClassB. The problems with classB traffic are related to its mid-level properties (TBD):
 - 1) ClassA1. Although similar to classA1, distinctions are needed to avoid:
 - i) Dual-queue. The STQ-size should not limit classB allocation levels.
 - ii) Single-queue: The absence of an STQ should not limit classB allocation levels.
 - 2) ClassC. Although similar to classC, distinctions are needed to ensure pre-emptive fairness.

9.2 ClassA0 dependencies

9.2.1 ClassA0 blockage scenario

The purpose of the secondary transit queue (STQ) is to queue incoming traffic until transient congestion conditions can be communicated and resolved. As such, partial filling of the STQs is not necessarily an uncommon occurrence, but something that happens during normal operations. As such, the allocation and fairness protocols should function correctly when the STQs are empty, slightly filled, or nearly full.

To illustrate potential STQ filling problems, consider the traffic loads of Figure 9.1a applied to topology of Figure 9.1b, leading to the classA0 bandwidth guarantee failure of Figure 9.1c. For this illustration, assume:

- a) Stations S1-to-S30 generate cumulative classC traffic loads of $.99 * linkRate$.
- b) Stations S30-to-S31 simultaneously burst subclassA0+classB traffic at 3% of link rate.
- c) Stations S1-to-S61 transmissions are destined for the S62 station.
- d) Stations S31-to-S60 have large STQs, with high&low thresholds of STQ/4 and STQ/8.
- e) Single-queue station S61 is allocated $.01 * linkRate$ classA0 capacity.
- f) The ring latencies are dominated by a long local round trip time (LLRTT) of the S30-to-S31 link

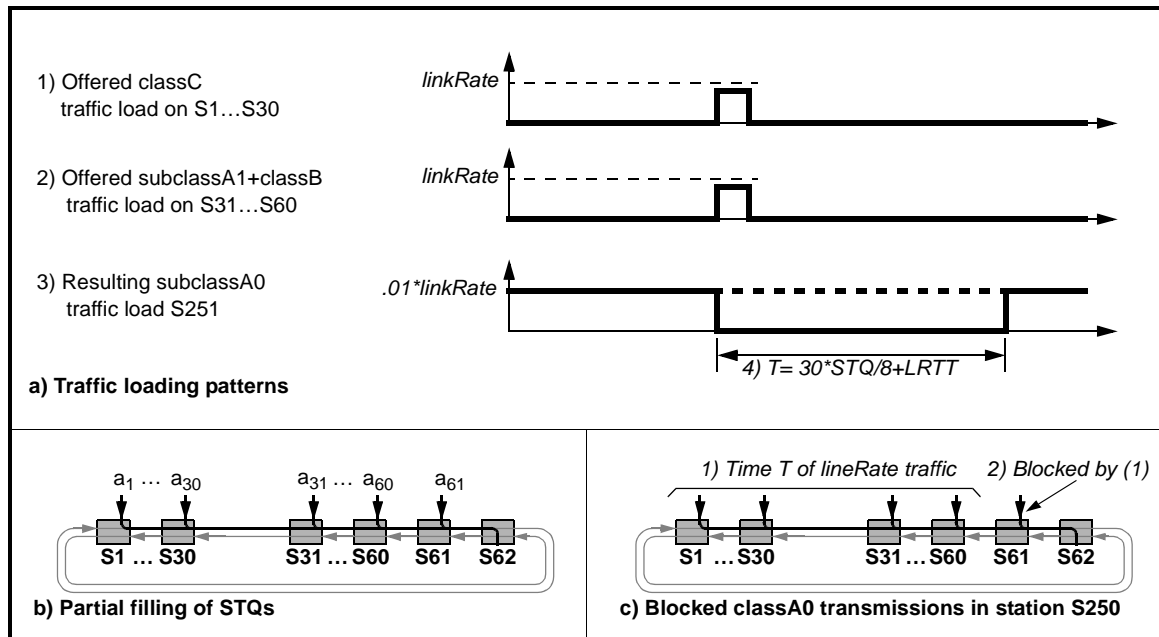


Figure 9.1—ClassA0 blockage scenario

With an unfortunate timing of the offered $a_{31} \dots a_{60}$ add traffic, their STQs can fill to the low threshold (1/8 of the STQ) before a congestion condition is communicated. An additional delay of LRTT occurs before the cumulative upstream classC traffic from $a_1 \dots a_{30}$ can be stopped. As a result, the cumulative fill levels of the S31...S60 queues equals a time duration $T = 30 * STQ / 8 + LRTT$.

There is no downstream shaper on the S31...S60 STQs, they continue to transmit at the *linkRate* until their STQs have emptied. Thus, the downstream S62 station is blocked for at least the duration T , which directly effects the worst case classA0 traffic jitter. The LLRTT delays are significant; for large buffer sizes, the STQ related delays become intolerable.

9.2.2 ClassA0 blockage avoidance

The classA0 blockage was caused by the partial filling of upstream STQ buffers, allowing the upstream stations to effectively exempt the transiting traffic from downstream shapers. Two types of solutions are therefore possible:

- a) Revise the transmission protocols, so that upstream STQs are never filled.
- b) Improve the effectiveness of the downstream shaper, so that STQ traffic is no longer exempted.

The viability of a type (a) solution is unlikely and validation proofs would be complex: the STQ are inherently intended to be partially filled under transient loading conditions. Attempting to restrict their filling, when detrimental effects are predicted or sensed, is likely to compromise their primary purpose of holding incoming opportunistic traffic while flow-control indications are being invoked.

The viability of a type (b) solution viable and validation proofs are manageable: the downstream shaper is applied to all conflicting traffic: STQ and added classB, in addition to added classC. There is no need to limit the added classA1 traffic, since subclassA0 and subclassA1 traffic have equal precedence.