9.3 ClassA congestion management

SubclassA0 support involves shaping the cumulative non-classA traffic to ensure sufficient downstream subclassA0 traffic, as illustrated in Figure 9.7. In this illustration, support of the single-queue downstream station S250 requires throttling of classB/classC in the upstream dual-queue stations.

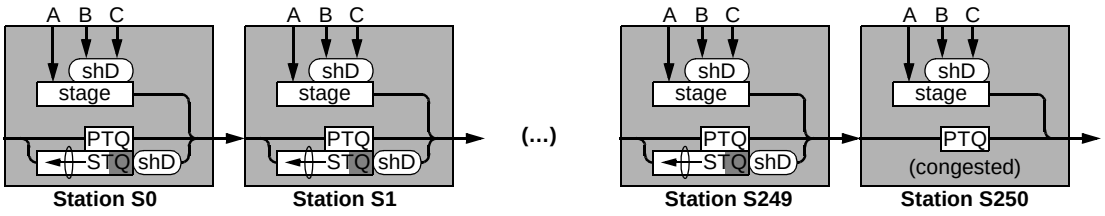

Figure 9.7—Queued upstream traffic

Within large topologies such as Figure 9.7, a modest 1ms level of queued STQ traffic in stations S0-to-S249 leads to a cumulative continuous traffic load of 1/4 second. To avoid unacceptable subclassA0 delays, station S250 transmissions are allowed to continue while upstream STQs are being drained. This is accomplished by shaping the upstream classB/classC add and transit transmissions, via one instance (with two traffic-policing locations) of the downstream shaper shD.

The downstream shaper is bimodal, in that its rate limit depends on the downstream congestion conditions. When only station S250 is congested, upstream rates are limited by the need to sustain prenegotiated levels of downstream subclassA0 traffic. If dual-queue station S249 were to become congested (as its STQ starts to fill), upstream dual-queue rates are limited to sustain cumulative (subclassA0 and subclassA1) levels of downstream classA traffic.

Throttling of upstream classB and classC traffic allows dual-queue stations (such as S249) to dedicate most of their surplus STQ space to the support of subclassA1 traffic, independent of the upstream classB loading conditions. The classB transmission opportunities are not lost: when classB throttling subsides, the accumulated classB-CIR credits are exercised with precedence over fairness-eligible transmission requests.

Editor DVJ's Notes: *To be removed prior to final publication.*
Specify how the classB traffic has precedence over classC traffic. This could affect the fairness protocols.