Stochastic Fair Queuing and Deficit Round Robin

Stochastic Fair Queuing (SFQ) [1] is a modification to the Fair Queuing scheme, that try to override its limitations.
With SFQ, a hash table is used to map packets to corresponding queues. Therefore, the number of queues may be considerably less than the number of possible flows. All flows that happen to hash to the same queue are handled in the same way, just as they were the same flow. Obviously, this approach strongly reduces the number of required queues. The disadvantage of this scheme is that flows that collide with other flows are treated unfairly. Therefore, the fairness guarantees are probabilistic; hence the name of Stochastic Fair Queuing. However, if the size of the hash index is sufficiently larger than the number of active flows, the probability of unfairness is small. Notice that, in order to have good performances, the number of queues needs only to be a small multiple of the number of active flows (which is much less than the number of possible flows, as required b Nagle's scheme).
A further development of the round-robin approach has led to the Deficit Round Robin scheme [17] where the fairness issue related to the presence of different packet sizes is addressed. In the DRR scheme SFQ is used to assign flows to queues. Queues are then served in round-robin order, and each queue is allowed to send a given amount of bytes (quantum) in each round of the robin. If a queue is not able to send a packet, because its size is too large, the remainder from the current quantum is added to the quantum for the next round.

Bibliography

[1] Golestani "A Self-Clocked fair queueing scheme for boradband applications" IEEE INFOCOM '94 pp 636-646 June 1994.