David Mulveney

"Core-Stateless Fair Queueing: Achieving Approximately Fair Bandwidth Allocations in High Speed Networks"

The key point of this paper is developing a new method for queueing packets in routers and making sure different flows only receive their fair share of the available bandwidth. This is accomplished by having edge routers in the Internet do some processing on each packet to put identifiers into the header giving a probability the packet should be dropped if it encounters a congested link. This probability is determined by the rate packets arrive and the available bandwidth in the link.

This approach performs well under most circumstances. CSFQ has a few problems when it comes to bursty traffic as shown in the ON-OFF traffic model experiment. Under this experiment CSFQ does not perform as well as FRED because it allows for a larger average queue size and therefore allows for more congestion in intermediate links. Another problem addressed in the paper is if edge routers put incorrect information into their packets. If the rate put in by a router is more than the actual rate CSFQ will drop that flows packets more often when they encounter congestion, but if the rate inserted into the packet is less than the actual rate the flow will get more bandwidth than it is supposed to. After some time CSFQ will realize that this misbehaving flow is sending more traffic than it says it is and CSFQ will eventually realize this and drop the flow's available bandwidth to correct the error. Based on this it is evident that CSFQ can be fooled to send extra traffic for short bursts, but it corrects the allocations after time.

There are some interesting clean slate ideas presented in this paper. It starts by using the idea of a bit stream abstraction for the traffic flowing through a link. The basis of the protocol is from this idea, using a round robin scheduler to ensure fair bandwidth allocation for all flows going through a link. For incremental deployment this bit stream abstraction is changed to a packet abstraction and the algorithm was then designed to give fair bandwidth allocations to each flow. There are different areas of the Internet that would benefit from deployment of this protocol. The routers on the edge would be able to perform the calculations needed faster and more efficiently than routers in the core. This would mean that edge routers should do most of the work so the routers in the core can read the information and route the packets to their destination quicker.