Steven Kappes

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

This paper presents another method of fair queueing by implementing max-min fairness. The network is divided into edge routers, which have fewer flows, and core backbone routers, which need to be able to operate at higher speeds. The edge networks calculate the rate of each flow. Every packet for a flow is then labeled with its rate. The core routers will then approximate max-min fairness by estimating the fair share rate of each flow. If there is congestion, packets will be probabilistically dropped to ensure flows operating above their fair share get scaled back.

The benefit of separating these operations is that the complexity is reduced at core routers. Core routers will not need to keep any per-flow state, only averaging the fair share rate, the aggregate arrival rate, the accepted traffic rate. This allows the complexity per packet to be only O(1) at core routers, while the edge routers should be able to handle classifying each packet into a flow. Simulations show that this approximation performs much better than non fair queueing algorithms such as FCFS and RED. CSFQ even outperforms deficit round-robin when there is a large number of flows.

Implementing such a scheme would certainly be more difficult than any of the other queueing methods, however. The network would need to be explicitly structured between core and edge routers. The scheme would need uniform deployment inside the island to implement CSFQ. If the island of routers is a single AS, this would not be a problem. A larger island would require co-operation though. If an edge router in the island did not label packets, core routers would not be able to approximate fair queueing. This ultimately breaks the end-to-end argument by adding functionality (packet labeling of a flows rate and evaluation) to the network. The paper also does not indicate how packet dropping will occur, if necessary, on edge routers. Presumably they could perform the same operations as core routers.