The Performance of Circuit Switching in the Internet

Pablo Molinero-Fernández, and Nick McKeown, Stanford University

{molinero,nickm}@klamath.stanford.edu

ABSTRACT:

Recently there has been renewed interest in circuit switching because of the ease of building very high capacity optical circuit switches. This poster studies the performance, as seen by the end user, of an Internet that uses circuit switching instead of, or in addition to, packet switching. TCP represents over 90% of current Internet traffic, which makes our usage of the network very connection oriented. Therefore, the key performance metric for the end user is the completion time for each connection, which we refer to as response time.

On the face of it, simply considering circuit switching as an alternative would seem a pointless exercise; the Internet is packet switched, and was deliberately built that way to enable the efficiencies afforded by statistical multiplexing, and thus it should provide a response time that is faster than that of circuit switching. However, link utilization is low, and falling, particularly at the core of the Internet, which means that statistical multiplexing is less important than it once was.

In this poster, we explore the performance of a network where a new circuit is created for each application flow, with particular emphasis on the response time experienced by users. We use simple M/G/1 and M/G/N queues to model application flows in both packet switched and circuit switched networks, as well as ns-2 simulations to validate the results. We conclude that because of high-bandwidth long-lived flows, it does not make sense to use circuit switching in shared access or local area networks. At the same time, our results suggest that in the core of the network, where access links limit the maximum flow rate, and where high capacity is needed most, there is little or no difference in performance between circuit switching and packet switching. Given that circuit switches can be built to have higher capacity than packet switches, this suggests that a circuit switched core warrants further investigation.

There are other performance metrics that are also very important when deciding whether circuit switching is a good technology for the Internet, such as scalability, simplicity, cost, reliability and availability, Quality of Service (QoS), or traffic engineering. We have used such criteria in other studies, but they will not be discussed in the poster with great detail given the space constraints. The end result is that we greatly benefit from the use of circuit switching in the backbone of the Internet.

Our current use of the network is very connection oriented, and, thus, it fits well with a network core that is circuit switched, while the edges use packet switching. We may just use a gateway that maps packet switched flows (e.g. TCP connections) to circuits. This gateway creates and destroys circuits as flows come and go. Then the deployment of circuit switching in the core does not require any change in existing end hosts and routers, which will continue to use IP the same way as today, as the gateways will hide the existence of circuits in the core. With this network architecture, circuit switched clouds can be inserted gradually, without requiring a flag day in the Internet.

FOR MORE INFORMATION:

`http://klamath.stanford.edu/TCPSwitching/`