Influence of Recovery Time on TCP Behaviour

Chris Develder, Didier Colle, Pim Van Heuven,
Steven Van den Berghe, Mario Pickavet, Piet Demeester
Ghent University - IMEC, Department of Information Technology
Sint-Pietersnieuwstraat 41, 9000 Gent (Belgium)
tel. no. +32 9 267 35 93 — fax. no. +32 9 267 35 99
e-mail {chris.develder, didier.colle, pim.vanheuven, svdberg, mario.pickavet, demeester}@intec.rug.ac.be

MPLS has been introduced a few years ago with as major motivation the potentially faster forwarding of IP packets than with classical IP routers. As time went by, the focus of MPLS has shifted somewhat towards traffic engineering. Yet another advantage of MPLS has recently emerged: MPLS allows the application of recovery schemes formerly confined to pure transmission networks. These recovery mechanisms enhance the bandwidth efficient IP networks with fast protection switching capabilities. Therefore, a better (faster) protection against network failures can be guaranteed.

The various MPLS recovery mechanisms as proposed in the IETF differ on several aspects, such as the speed of recovery, the additional network capacity requirements they impose, the cost in terms of extra control traffic and processing power, and the failure scenarios against which they protect. It is quite clear that trade-offs will have to be made, e.g. between cost and recovery speed. To help solving these kinds of dilemmas, studies are needed to quantify the pros and cons of the different mechanisms.

This paper focuses on the speed of the recovery mechanisms for a very particular traffic type. Based on a thorough study of simulation results, we investigate the importance or potential nuisance of very fast protection switching for TCP flows. When a protection switch is carried out within a best effort network (e.g. as illustrated in Figure 1), the protection switch will affect the flows being switched to an alternative path (the so-called switched flows). In addition, the so-called fixed flows, i.e. those already traversing the links of the recovery path before the switch occurred, will be influenced. Therefore, we investigate the effects of protection switching on both flow categories. We study the impact of the time it takes to recover from a failure, on the goodput attained by the different flows.

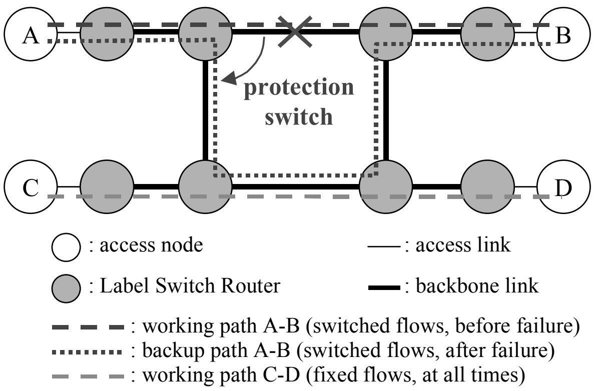

Figure 1. Network topology and paths followed by TCP flows. The so-called switched TCP flows will be set-up between A and B, the fixed ones between C and D.

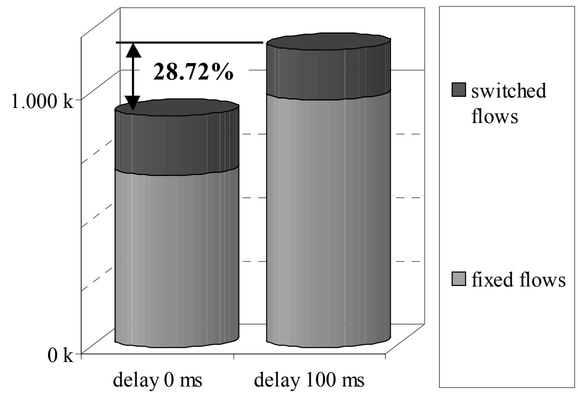

Figure 2. Total goodput (in bytes) during the first second after link failure for both fixed and switched flows for different protection switch delays, using TCP New Reno. (Note: RTT is 100-150 ms.)

From detailed simulations, we infer that very fast protection switching for TCP flows is not a must. On the contrary, a certain protection switch delay is preferred over an immediate protection switch: it allows for a greater amount of traffic to be successfully transported end to end, within the same time interval — this is what end-users will care about. In Figure 2, one can see that by introducing a protection switch delay, the switched flows lose a little, but the fixed flows' goodput increases, resulting in a net gain of 28.72% for this particular case (using the topology depicted in Figure 1). This possibly somewhat surprising result stems from the dynamic behaviour of TCP.

ACKNOWLEDGEMENTS
Part of this work has been supported by the European Commission through the IST projects LION and DAVID and by the Flemish Government through the IWT project ITA/980272/INTEC, an IWT scholarship, and an FWO scholarship.