Design and analysis of opticalpacket switching networks

Ontwerp en analyse van optische pakketschakelende netwerken

Chris Develder

Despite the recent economic fallback, following the explosion of the technology bubble, we still observe a steadily continuing growth of telecommunication traffic. It's no secret that the originally dominant voice traffic nowadays is dominated in volume by data traffic, especially originating from the Internet. That Internet, being characterized by packet based transport of heterogeneous information, has inspired people to introduce the idea of applying the packet based switching concept also in the underlying transport networks. This resulted in the proposition of adopting an Optical Packet Switching (OPS) network. This PhD dissertation aims at identifying the major issues and difficulties that arise in OPS networks, from a logical performance point of view. The latter means that we will not deal so much with physical/technological challenges of OPS, but rather discuss how to handle the optical packets in terms of contention resolution, routing, etc.

To the field of OPS architectures we contributed with two studies. A first focused on node architectures for metro area networks (MANs). The metro environment is seen as the place where OPS may be introduced on a relatively short term. To compete with recent alternatives such as E-PONs or RPR, the cost of an OPS network should be kept low. To allow for relatively simple node architectures, various projects have proposed ring-based MANs. We compared two node architectures in terms of the amount of required resources to set up a given traffic matrix for a given set of nodes. We showed that a simple architecture without any active switching components in the transit path (followed by traffic passing through a MAN ring node) does not lead to any penalty in terms of transmitter and receiver capacity. For a more distant future, we addressed the question how to scale OPS switching nodes for a meshed backbone network. Our original contribution showed how a Clos-like approach, with two switching stages and a wavelength conversion stage, allows significant reduction of the number of components in a SOA-based broadcast-and-select switching matrix.

Next, we move on to the problem of scheduling packets in such an OPS switch. We considered a generic architecture with a shared feed-back buffer, comprising fibre delay lines (FDLs). In this packet switch we assume contention will be resolved using wavelength conversion and recirculation through the FDL buffer. For the case of fixed length packets and a slotted operation of the switch, we introduced a new Balance algorithm which leads to far lower packet losses than previous proposals. In case of variable length packets, we have assessed how they can be dealt with in a slotted OPS switch: how do we schedule a train of slots which comprise a single variable length packet? We considered a train approach, making a scheduling decision for the complete train at once, upon arrival of the first "wagon". This train approach was compared with a wagon approach where the switch schedules each slot independently. We found that only for a particular range of loads, the wagon approach may attain somewhat lower loss. Yet, with a simple priority mechanism, the wagon approach can achieve quite strong service differentiation, which is not the case for the train approach. For the case of switching variable packets in an asynchronous (i.e. unslotted) packet regime, we assessed various alternatives to offering QoS. We introduced two new approaches: a look-ahead approach, and a slotted control approach. We found that the slotted algorithm, despite its simplicity, may attain acceptable performance levels if the load is relatively low. Still, far stronger differentiation is reached with the look-ahead approach. We obtain similar or slightly lower loss rates compared to the well-known approach of offset differentiation (cf. OBS), but more strongly discriminate against longer packets.

On a network-level, we developed a routing algorithm aiming at minimizing the maximal loss rate occurring in a given network carrying a given set of traffic flows. We showed that through careful estimation of the loss rate, we can reach significantly lower loss levels, not only compared to a simple shortest path routing, but even compared to a load balancing algorithm.

Also the applications making use of an optical transport network were considered. We provided an extensive original study of the effects of protection switching actions on TCP performance. We showed the detrimental effects of interaction between switched flows and those already present on (part of) the protection route. Also, we quantified the effects of changing path lengths and the influence of the speed of protection on TCP goodput. We culminated the gathered knowledge in a case study comparing different MPLS-based protection schemes.