Seminar Summary: Design and Implementation of Scheduled Optical Flow Switching

Bishwaroop Ganguly

November 12, 2001

Abstract

This document outlines a discussion of Scheduled Optical Flow Switching (OFS), an ongoing WDM optical network architecture study. In this context OFS involves on-demand creation of optical connections for data transport between end-hosts. The challenges arise from the fact that these connections are of short duration (less than one second), and that current optical network technology supporting such an infrastructure is in its infancy. In this seminar, we motivate the overall idea of OFS. We then briefly describe analysis and modeling that motivates our particular architecture choices. Finally we discuss and present results from an actual OFS implementation we have done on the ONRAMP optical access area network. Our results show that 90% efficiency is possible for one-second optical connections, with a minimum of specialized hardware or software support.

OFS Motivation

Current Wide-area or Access-area optical data networks are generally statically or slowly reconfigurable. They usually consist of large data trunks, terminated at each hop by an electronic router. Reconfiguration, if any, is usually achieved on the order of hours if not days. In contrast to this, Optical Flow Switching is an idea that involves on-demand dynamic creation of all optical paths from user to user in much less than a second.

The primary benefit of this type of OFS is *router bypass*. Clearly these on-demand all-optical transactions will not need to be switched by electronic routers or even converted from an optical to electronic signal. This has the potential of a large cost savings for network builders, as well as increasing the *effective capacity* of the routers. A further benefit is transparency of these user to user optical connections, which can be used to provide whatever Quality of Service (QoS) is required for the application.

A number of new optical device technologies are enabling this. Among these are faster optical switch and tunable transceiver technology. An open question is that of: what is the best architecture in which to deploy these devices and what are the potential benefits?

Our goal is to investigate the use and feasibility of Optical Flow Switching (OFS) in the context of an access area (metropolitan area) network. We have implemented OFS in the ONRAMP network with a goal of demonstrating viability of our approach as well as efficiency/performance. We also want to uncover any stumbling blocks this approach has as the network scales to even larger sizes.

Scheduled OFS

The important design decision that we have made in OFS is: use of timing information in connection setup. Almost all other approaches described in the literature use either link-state updates or probing of the network to acquire information about availability of optical resources. This includes the currently discussed MP-lambda-S suggested by the corresponding working group of the IETF. The latency involved in these approaches cause the information obtained to be *stale*. This can have detrimental effects on network performance as we will show.

We instead use timing in the network to maintain accurate information about the future state of flows in the network. This allows high utilization of the network resources, perhaps at a cost of small delay to the user. In this presentation, we discuss using a time-slotted approach to connection request and setup. However, we emphasize that there are many other approaches that can be used with timing information, including some that allow connection setup based on perfect network state information.

OFS Demonstration and Results

We have designed and implemented scheduled OFS in the ONRAMP optical access area network. ONRAMP stands for Optical Network in the Access Area using Multi-wavelength Protocols. It is an access area network deployed in the Boston metropolitan area.

Scheduled OFS on this network uses network-global timing and synchronization. This implementation involves a number of issues ranging from O/S support to optical device optimization. We will discuss these in brief. We will also discuss transport layer issues involved in OFS focusing on a design decision between a connection-based vs. a connection-less transport layer.

Our experiment involves running OFS from one user to another user over ONRAMP, switching wavelength and route at each timeslot for a variety of timeslot sizes. Since reconfiguration of the network is required to achieve this, we expect some loss of efficiency versus a static optical connection approach. Our goal is to minimize this overhead, and examine what the key issues are.

Performance results for OFS are measured in terms of user-to-user throughput of IP packets compared to maximum possible throughput for a static optical connection. We find that we achieve 90% efficiency for one-second timeslot transactions, multiplexed over a number of wavelengths and over both possible routes in ONRAMP. For these results we did not use any specialized OS or hardware support. We will discuss what the limiting factors were for even better performance, and show that these can be addressed with a modest amount of hardware support.