Optical Crossconnect with Shared Wavelength Conversion under Dynamic Loading

Debasis Mitra, Carl Nuzman, Iraj Saniee, Phil Whiting
{mitra,nuzman,iis,pwhiting}@research.bell-labs.com
Bell Laboratories, Lucent Technologies
600-700 Mountain Ave, Murray Hill, NJ 07974

Abstract *We examine the blocking performance of all-optical crossconnects with a limited ability to convert signals between wavelengths, under a dynamic traffic model. In the model, traffic requests arrive in Poisson streams, and accepted lightpaths are held for random holding times. A relatively small regenerator pool is sufficient for high switch utilization*

Index terms - crossconnect, wavelength conversion, transparency, traffic models

1 Introduction

Crossconnects in today's optical core networks rely on optical-electronic-optical conversion and high-speed electronics. As the rates and volumes of data traffic in optical core networks scale up, the high costs, space requirements, and power requirements of electronic crossconnects become prohibitive, and all-optical crossconnects provide a scalable alternative.

In the near term, all traffic passing through an all-optical crossconnect will be received and regenerated, using a regenerator on every port of the crossconnect. This allows traffic engineering to be done on a link-by-link basis, and allows signals to be packed efficiently without wavelength blocking. On the other hand, as long-reach line systems are deployed, many network demands could potentially be routed end-to-end without regeneration. With sufficiently long reach or a sufficiently local network, no demands would require regeneration, and crossconnects could be provisioned without regenerators at great cost savings. Such a network would also be transparent to signal format and bit rate. However, a network without wavelength conversion may suffer significant wavelength blocking at the nodes, so that the network would need to be designed with a large amount of spare capacity. In this paper we consider provisioning an all-optical crossconnect with a shared pool of regenerators. This allows intermediate points in the tradeoff between the cost of regeneration and the cost of spare capacity. Such a crossconnect is *selectively transparent*. The concept of selective transparency is explored for a network with fixed demands in [1]. Here, we focus instead on a single node, under a dynamic traffic model.

2 Crossconnect and Traffic Model

We consider a single node in an optical wavelength division multiplexed network. The node has D neighbors, with M fibers to and from each neighboring node, and with each fiber carrying W wavelength channels. (The symmetry of the node is for simplicity of exposition, and is not an essential characteristic of the model.) The signals are demultiplexed and fed into and all-optical switch fabric, where $N >= MDW$. There is also pool of C regenerators which can be shared arbitrarily by the lightpaths passing through the node. This pool could be implemented, for example, by connecting the regenerators in a loopback configuration leading from a set of switch output ports back to a set of switch input ports. In this case, an NxN optical switch fabric would be required with $N \geq MDW$. In our analysis, we consider regenerators which can receive and transmit signals of any wavlength. By simulation, we can also consider regenerators with fixed output wavelength.

In our traffic model, connection requests arrive on each of the MDW input channels of the crossconnect in Poisson streams. Each connections request specifies a a particular neighboring node to which the connection should be routed. If it is possible for the connection to be carried, it is accepted; otherwise the connection is blocked. Accepted connections exists for a random holding time before being taken down. The traffic rate ν on each channel is measured in units of *erlangs*, given by the mean number of arrivals during an average holding time. For simplicity of exposition, we assume all input channels share a common rate ν, and that each connection is equally likely to be routed to any of the neighboring nodes. Given the dimensions M,D,W of the switch, the size C of the converter pool, and the offered traffic rate ν, our first goal is to determine the connection blocking probability.

The following algorithm is used to determine whether or not a given connection can be carried, and if it is carried, which resources (fibers and converters) are used. A particular connection request arrives on a channel with wavelength $1 \leq w \leq W$, destined for neighboring node $1 \leq d \leq D$.

- If any of the M channels with wavelength w leading to node d is available, accept the connection. Choose one of these channels at random.

- Otherwise, if a regnerator is available, and if there are any channels available on the fibers leading to d, accept the connection. Use a regnerator and choose one of the available channels at random.

- Otherwise, reject the connection.

3 Blocking and Capacity

Under the dynamic traffic model with $\nu > 0$ and $D > 1$, it is always possible that more requests will desire a particular destination d than can be carried on the M fibers leading to it. To define the capacity of the crossconnect, we can set a maximum tolerable blocking probability $p_{\max}$, and determine the largest traffic rate ν which satifies the blocking constraint.

3.1 Analysis

When the number of regenerators C is equal to the maximum number of connections that can be carried MDW, blocking only occurs when the set of fibers leading to a given destination are completely full. The fibers leading to the destination form an Erlang loss model with MW resources and total Erlang load νMW. Hence the blocking probability is given by Erlang's formula, denoted $E\,(\,\nu MW\,;\,MW\,)$.

When there are no regenerators, each wavelength forms a separate plane in the switch. In this case, the M channels of a particular wavelength leading to a particular destination are the resources, the load is νM, and the blocking probability is $E\,(\,\nu M\,;\,M\,)$.

When the number of regenerators lies between these extremes, a closed form for the blocking probability is not available. However, an analysis technique based on fixed point iteration provides an accurate approximation. [2].

3.2 Discussion

Figure 1 shows the blocking behavior of a node as a function of the Erlang load per channel. This node has $D = 5$ neighbors, $M = 4$ fibers per link, and $W = 30$ wavelengths per fiber. The curves are given by analytical methods, while the circles are data points produced by Monte Carlo simulation. When $C = 0$, the wavelength contention strongly limits the traffic that can be carried, to less than 10% of the total switch capacity for a blocking probability of 10^{-4}. With full wavelength conversion, the utilization increases dramatically to 70%. Using a regenerator pool of just 60 of a possible 600, the utilization is a respectable 50%. Here the tradeoff between the number of converters and the spare capacity is apparent.

To gain further insight, we can examine the way in which the blocking probability falls as C increases, as in Figure 2. On the left side of the figure, the regenerator pool is small, and the blocking is dominated by wavelength contention. On the right side of the figure, the regenerator pool is no longer a constraint, and the blocking approaches that of a switch with full conversion. There is a natural pool size C^*, in this case around $C^* = MDW/2$, to which the pool can be reduced without significantly increasing the blocking probability. In many cases, particularly when $M > 1$, the natural pool size is much smaller than MDW.

Although connection-based models such as the one considered here are used with great success and accuracy in the study of telephone networks, they may not provide a quantitatively accurate model for lightpath demands in optical core networks. Nevertheless, the results of this study suggest that dynamism in the network can lead to significant wavelength blocking, which may be mitigated by small but flexible wavelength conversion resources. Sources of dynamism in future optical mesh networks may include long term demand fluctuations as well as shorter time scale changes driven by network impairments or bandwidth trading.

References

[1] R. Giles, K. Kumaran, D. Mitra, C. Nuzman, I. Saniee, and P. Whiting, "Selective Transparency in Optical Networks", submitted to OFC 2002.

[2] D. Mitra, C. Nuzman, I. Saniee, and P. Whiting, "An Analytical Model for Connection Blocking in Optical Crossconnects", in prepration.

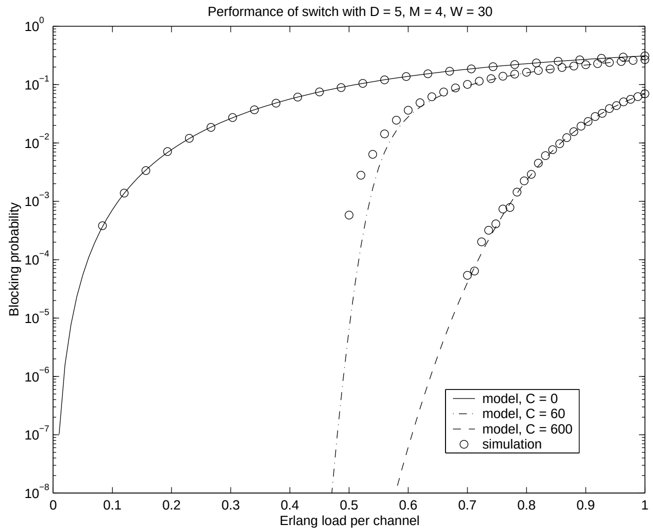

Figure 1: Blocking Probability as a Function of Load, for a Crossconnect with Shared Regeneration

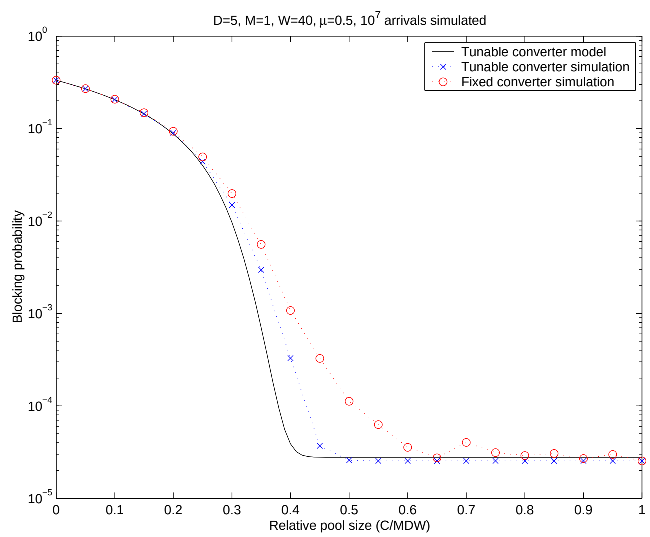

Figure 2: Blocking Probability as a Function of Pool Size, for a Crossconnect with Shared Regeneration