

The Effect of Bridge-and-Roll on Minimizing Wavelength Conversion for Dynamic Traffic

Steven Fortune(1), Lisa Zhang (2)

1,2 : Bell Labs, Alcatel-Lucent. 600 Mountain Ave, Murray Hill, NJ 07974 USA. {sjf, ylz}@bell-labs.com

Abstract

The number of wavelength conversions required to handle dynamic traffic in a WDM network is significantly more than for comparable static traffic. However, this excess can be substantially reduced with bridge-and-roll.

Introduction

Increasing optical transparency is a paramount goal for modern WDM networks. Clever wavelength assignment helps reduce wavelength conversion and hence increases transparency. Wavelength assignment for static traffic has been extensively studied (e.g. [1,2] and references within), with some attention also given to dynamic traffic (e.g. [3] studies schemes to reduce blocking given constrained wavelength conversion). In this note we quantify the excess conversions required for dynamic traffic over static traffic, and observe that bridge-and-roll significantly reduces this excess.

In the static case, a network topology and a demand matrix are fixed. Each demand is provisioned with its route, i.e. the sequence of fiber links connecting the demand source to destination, and its wavelength assignment on each link. A wavelength converter is required whenever the wavelength changes along a route.

In the dynamic case, demands arrive and depart in time. When a demand arrives, it must be provisioned from available wavelengths without affecting other demands; when the demand departs, its wavelengths are freed.

In the dynamic case with *bridge-and-roll* (BR), some number of existing demands may be reprovisioned along with the new demand. The reprovisioned demands may only use available wavelengths.

Methodology

We wish to compare the number of converters required for static and dynamic traffic, with and without bridge-and-roll. However, wavelength assignment and routing is NP hard and hence cannot reasonably be solved to optimality. To bound the number of variations in the comparison, and to focus explicitly on wavelength assignment and not on routing, we proceed as follows.

We require that each demand follow the same route in the dynamic solution as in the static solution (wavelength assignment may differ). We also use heuristics for both routing and wavelength assignment.

In the dynamic setting, the route for a new demand d is found using a shortest-path algorithm that is

aware of available capacity but not wavelength assignment. Once the route for demand d is found, we first choose a wavelength that traverses as far as possible along d 's route until a blockage occurs. We then repeat until every link on the route is assigned a wavelength. This greedy approach is locally optimal, i.e. requires the fewest converters if demand d is considered in isolation.

For a fixed dynamic scenario, i.e. network and demand schedule, we let $D(t)$ be the number of converters required at time t .

For bridge-and-roll, we choose to reprovision all existing demands that currently require wavelength converters. The wavelength assignment along a demand route may change, but the route itself cannot. We let $R(t)$ be the number of converters required at time t with bridge-and-roll.

To simulate static wavelength assignment, for each time t we reassign wavelengths for all active demands using the routes found in the dynamic case. This wavelength assignment is performed incrementally using the local heuristic. We let $S(t)$ be the number of converters required.

We now use

$$D = \max_t D(t), S = \max_t S(t), R = \max_t R(t)$$

to compare converter costs among the three cases.

Even for a fixed dynamic scenario, the actual values of D , S , and R are dependent upon many choices. For example, for shortest-path routing we can choose many edge weights (unit length, physical length, or a function of residual capacity). Similarly the incremental wavelength assignment depends upon the order in which demands are considered (e.g. random, longer routes first, highly congested routes first, routes previously requiring many converters first). We believe that the comparative behaviour of D , S , and R does not depend strongly upon which choice is made, as long as the choices are the same in each case; plenty of experimentation is possible.

We used six data sets that are similar to those from major carriers in the US and Europe. See Table 1. Each data set provides a network topology and a static traffic matrix M . Network sizes vary from roughly 11 nodes and 14 links to 51 nodes and 72

links; degree 2 nodes have been eliminated. In each case we created a dynamic demand scenario from the traffic matrix M as follows. Between each node pair (s,t) , we generate a stream of demands in a Poisson process such that the ratio of the expected holding time to the expected interarrival time is a constant times $M[s,t]$ and the constant is chosen as large as possible so that no demand routing ever failed. We also generate non-Poisson traffic by having deterministic holding and interarrival times or choosing them from a uniform distribution.

Observations

For this traffic model, the peak number of demands is 15-30% more than the average number of demands (table 1).

Dynamic traffic (without bridge-and-roll) requires significantly more converters than static traffic. One cause appears to be that wavelength fragmentation can propagate in time. For historical reasons, it is possible that an old demand d requires a converter; suppose d starts on wavelength v and then switches to w . Suppose a new demand d' traverses the same route as d . At a time of heavy load, it may be that only wavelength w is available for the first part of the route, and wavelength v for the second. Hence d' is forced to use a converter as well. A static solution could eliminate both converters.

With bridge-and-roll, there is an opportunity for the fragmentation of the route of d to be eliminated before d' arrives. This effect is clear in Table 1, where bridge-and-roll substantially decreases the number of required converters. See also Figure 1, which shows a trace of the converter count as a function of time. At each time step the number of reprovisioned demands is typically fewer than 20% of the number of active demands, and is typically at most twice the number of new demands.

The order under which demands are assigned wavelengths is important. As in the case of static traffic, orderings that are based on some kind of priority (length, congestion, previous converter count) consistently outperform the ordering given by the input data, typically by 5-15%.

The scenarios described above have assumed that optical impairments do not limit optical reach. If the route length exceeds optical reach, an optical transponder (OT) can serve the dual purpose of regenerating an optical signal and changing wavelength. We reperformed several of the scenarios with reach limited to 60-80% of the average route length; the local wavelength assignment algorithm extends easily to this case. The total number of OTs increased; however the number of OTs required purely for wavelength assignment was 65-85% less than recorded in Table 1. Bridge-and-roll was still as advantageous.

Conclusion

Over several scenarios, dynamic traffic requires many more wavelength converters than comparable static traffic; however, bridge-and-roll can effectively reduce the number of required converters. This suggests that flexible implementation of bridge-and-roll within WDM switching elements can reduce converter costs. A detailed prediction of converter savings likely requires considerable exploration of the space of routing, wavelength assignment, and bridge-and-roll strategies.




Figure 1: Number of converters per time step. Top curve is dynamic traffic without bridge-and-roll $D(t)$, middle with bridge-and-roll $R(t)$, bottom with static traffic $S(t)$.

Instances	A	B	C	D	E	F
nodes	11	11	11	23	51	51
links	14	14	14	30	72	72
average utilization	35%	40%	42%	61%	48%	40%
max demands	51	55	56	332	587	246
average demands	39	44	47	271	514	206
D	28	44	54	48	185	107
R	13	19	17	39	126	89
S	6	10	11	20	72	71

Table 1: Data sets. Average utilization is the average percentage of occupied wavelengths per fiber. D is the maximum number of converters for dynamic traffic; R with bridge-and-roll; and S for static traffic.

References

1. Koster et al, Minimum Converter Wavelength Assignment in All-Optical Networks, ZIB-Report 03-45, 2003
2. Farrell et al, Wavelength Assignment in Optical Network Design. IMA-Report, 2007.
3. Spath et al, ECOC '98, pp 359-360.