ROADM Enabled Optimization in WDM Rings

Pankaj Risbood, Carl Nuzman, Nachi Nithi, Sanjay Patel
Bell Laboratories, Lucent Technologies
{risbood,cnuzman,nk2,sanjay}@bell-labs.com

Abstract. Deployment of reconfigurable OADMs in WDM rings is expected to bring large operational savings. Such nodes also enable online network optimization; we quantify the potential savings as a function of element functionality and traffic churn.
OCIS Codes: (060.4250) Networks; (060.2360) Fiber optics links and subsystems

I. Introduction

Transport networks have been traditionally designed for voice traffic whose behavior is predictable and has been well studied. With the rise of the Internet, traffic has become more data-centric. In-fact the volume of data traffic has already surpassed the voice and is growing at a much faster pace. Data traffic is burstier and less predictable than voice: in the future new applications like video-on-demand and online gaming, which require short-lived high bandwidth connections, are expected to create a need for on-demand wavelength services. If realized this trend will create churn even at wavelength level and is expected to drive the deployment of transport networks designed for flexibility.

Dense Wavelength Division Multiplexing (DWDM) has established itself as the cheapest and most reliable mechanism for transporting information bits in long haul as well as in regional metro networks. For reasons of resiliency and simplicity, DWDM networks are predominantly configured as rings, consisting of fiber links connecting Optical Add-Drop Multiplexers (OADM). In recent years Reconfigurable OADM (ROADM) architectures have been proposed in order to automate network management and allow the transport network to adapt more readily to changing loads [1]. In this paper we make a case that ROADM deployment not only reduces operating expenditure but the added flexibility helps save/defer capital investment. In particular we show that under churn a scheme for reconfiguring circuits carrying live traffic reduces request rejection and improves network utilization. We summarize extensive simulation results for various ring sizes, traffic patterns, loads and traffic churn.

II. Network Element Model

In this section we briefly describe various NE models that are used in experiments. Note that we only emphasize functionality; the actual technology used to implement the functionality is beyond the scope of this paper. A traditional OADM consists of a Mux-Demux pair and a set of add/drop ports for east and west side of transmission. A full system (except OTUs) needs to be deployed on day one (high startup cost). ROADMs offer a *banded* model minimizing one time capital expenditure (pay-as-you-grow) [2].
The second property we consider is wavelength and route flexibility, for which a wide variety of designs have been proposed [1,3]. In the simplest ROADM, each add-drop port is tied to a specific wavelength and ring direction. More advanced devices, which we refer to as *wavelength-flexible* ROADMs, allow each port to access any of a set of wavelengths. As in [4], we assume the wavelengths are divided into disjoint bands, and each port is assigned to one such band. The number of wavelengths in each band is called the band size. *Route-flexible* ROADMs allow a port to access either of two ring directions.

A wavelength and route flexible ROADM can be used to move a circuit, i.e. change its wavelength and/or route. In principle, the move can be made in a hitless fashion, similar to automatic protection switching. Combined with an effective network management system, such ROADMs would sufficiently reduce the operational costs and risks to enable operations that are infeasible in manually-controlled networks. In this paper, we study the degree to which the operation of online repacking of existing services can make room for future services, increase network utilization, and delay capacity exhaustion.

III. Capital savings

As stated earlier, a flexible ROADM controlled by an effective network management system enables optimizations like ring defragmentation [5], saving/deferring future capital expenditure. In this section we study effectiveness of such optimizations by means of extensive simulations. In our model, a service is a bi-directional communication channel between two network nodes. When a service request arrives, an attempt is made to assign it to a circuit, consisting of a route and wavelength. If no circuit is available, the ring is repacked. If there is still no circuit available after repacking, the service is blocked. A successfully assigned service may later depart, at which time the assigned wavelength becomes available again. The *first block capacity* of a network is defined as the number of services accepted before the first service is blocked.

We measured the first block capacity of various ring networks by simulation. In each scenario, the arrival and departure sequences of services were modeled using a Poisson arrival process and exponentially distributed service times. The node pairs for each service were chosen uniformly at random. The degree of churn is quantified by the Erlang load μ, which is defined as the product of service arrival rate and holding time. Infinite Erlang load (say infinite service times) corresponds to no churn, and the churn increases as μ decreases. In the results reported below, we examine the case of no churn and the case μ=4W, representing a high-level of churn in which departures are almost as common as arrivals on a fully loaded ring.

An arriving service is always assigned the shortest path (East or West) and the available wavelength with the lowest index. The repacking algorithm is a based on the method described in [5] in the context of SONET rings. That is, a *gap* is defined to be a consecutive sequence of unoccupied links on a given wavelength. A gap of m links is given the weight 2^{m-1}, and the packing metric for a given ring is defined to be the sum of the weights of all gaps. Repacking is done in a greedy sequential fashion; in each step the circuit move that causes the greatest increase in the packing metric is selected, until no single move is beneficial. The set of allowable moves depends on the flexibility and band size B of the ROADMs: if the ROADMs are wavelength flexible, a circuit may move to any other wavelength in its band, and if the ROADMs are route flexible, each circuit may toggle between East and West.

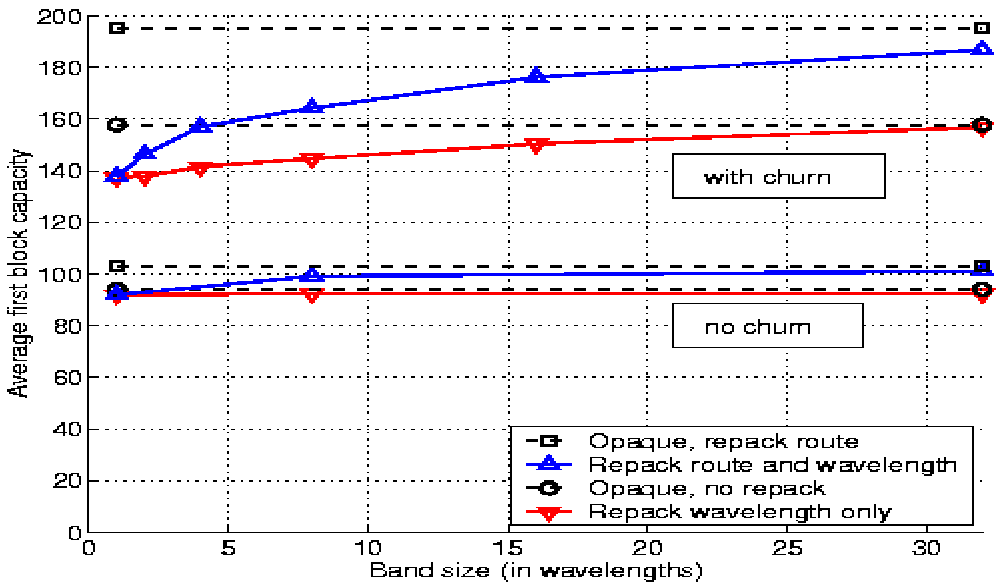

Figure 1: First-block capacity of 32 channel ring networks as a function of band

In order to upper bound the best that could be achieved with wavelength packing, we also considered *opaque* networks, in which a service is free to use different wavelengths on each link. To bound the best that can be achieved by repacking wavelength and route, we considered opaque networks with routes determined by the optimal ring loading algorithm [6]. Although we have studied a variety of ring sizes, the results presented are all for rings with eight nodes. The qualitative and quantitative results that we report are very insensitive to ring size, except for very small rings with say 3 or 4 nodes.

The effects of the above parameters on first block capacity are depicted in Figure 1, for an 8-node ring with 32 wavelengths. The ordinate of each data point in the plot is obtained by averaging the first block capacity over 500 random scenarios of service arrivals and departures, while the abscissa represents the band size B. The lower set of curves were obtained with no churn (i.e. no departures), and the upper set of curves were obtained with high churn (μ=4W). The red and blue curves depict the capacity as a function of B when repacking wavelengths or repacking wavelength and route, respectively. Each has a horizontal upper bound, representing the capacity of an opaque ring and that of optimal ring loading, respectively. Although the capacity with churn appears much higher than the capacity without churn, this is not a fair comparison since the former includes services that have arrived and departed before first block. However, comparing the case B=1 to B=W, it is clear that the relative increase in capacity due to repacking is much more significant with churn than without. Indeed, when there is no churn, the simple first-fit wavelength assignment algorithm packs the ring very effectively. Modifying routes without repacking wavelengths (blue curve, B=1) provides almost no benefit, repacking wavelengths without changing the route provides some benefit (red curve, B=32), but the greatest improvement is achieved by wavelength and route flexible ROADMs (blue curve, B=32). The largest incremental improvements in capacity with band size occur when the band size is small. For example, a ROADM with bands of size eight and the ability to reroute circuits performs better than a completely colorless ROADM that cannot reroute; it also performs better than a opaque network (full wavelength conversion) that cannot reroute.

Under our random service arrival model, the first block capacity is a random variable with considerable variance. For example, Figure 2(a) depicts the distribution of first block capacity corresponding to two data points with churn from Figure 1, for a network with and without wavelength and route repacking. Although the mean of the repack enabled distribution is 36% higher than that of the one without, there is significant overlap between the distributions. The time and expense of repacking operations could depend on the number times the ring is repacked and on the total number of circuits moved. Figure 2(b) depicts the distribution of the total number of circuits moved, for a network capable of repacking route and wavelength. The histograms with and without churn have mean values of 42 and 16, respectively.

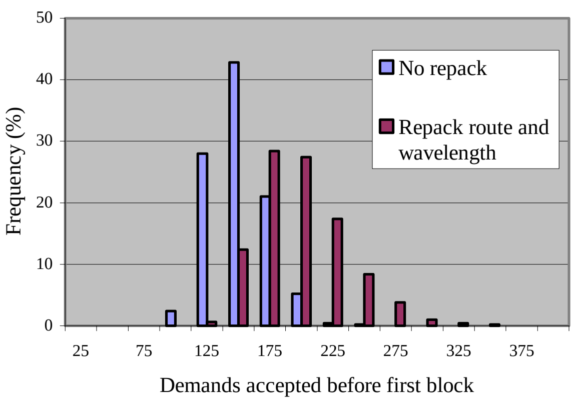

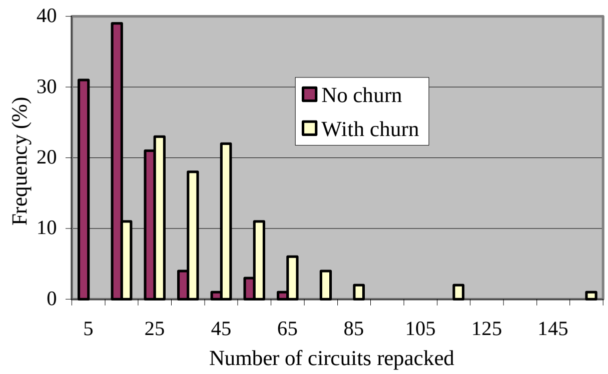

Figure 2 (a) Histograms of first block capacity with and without repacking (b) Histograms of the number of circuits repacked before first block with and without churn, for an eight node ring with 32 wavelengths.

IV. Conclusion

Repacking circuits on a WDM ring can significantly increase the first block capacity when there is significant service churn, while without churn the effect is marginal. A heuristic repacking algorithm based on exponentially weighted gaps is simple and effective; it provides a sequence of bridge-and-roll moves that can be performed online. Finally, adding routing flexibility to wavelength flexible ROADMs has a large impact on repacking effectiveness.

References

[1] A. Saleh et al, “Architectural principles of optical regional and metropolitan access networks,” J. Lightwave Tech. **17**, 2431-2448 (1999).
[2] C. Nuzman et al, “Effects of modularity and connectivity on OADM deployment in ring networks,” in OSA TOPS **86**, OFC Conference (Optical Society of America, Washington, DC, 2003), pp. WH-05.
[3] M. Feuer et al, “Routing power: a metric for reconfigurable wavelength add/drop multiplexers,” in OSA TOPS, OFC Conference (Optical Society of America, Washington, DC, 2002), pp. TuX-01.
[4] D. Cavendish et al, “Routing and wavelength assignment in WDM rings with heterogeneous wavelength conversion capabilities” in Proc. INFOCOM, (IEEE, New York, 2002), pp. 1415-1424.
[5] S. Acharya et al, “Hitless network engineering of SONET rings” in Proc. GLOBECOM (IEEE, New York, 2003).
[6] A. Schrijver et al "The ring loading problem," SIAM Review 41, 777-791 (1999).