Traffic Patterns and Network Churn

K.J. Warbrick

Nortel Networks

Abstract: Network analysts continue to report [1] that internet traffic continues to grow at about 100% per annum. This paper examines what network traffic conditions are necessary for the redeployment of wavelengths in the core of the network. The analysis, based upon a simple network model, allows the network value of wavelength reconfiguration to be explored.

1 Introduction.

With the continuing rise in offered traffic volumes to the core networks being generated by evolving services, increasingly using Fibre Channel and Gigabit Ethernet interfaces, together with the slowdown in development and deployment of the next generation line systems at STM-256 / OC 768 (~40 Gbit/s) there is the potential for the traffic aggregation balance to change.

Traditionally the telephone service traffic was aggregated through many levels from the 64 kbits/s voice channels through E1 channel (x30) through E2, E3 and E4, each of which aggregated by a factor of 4. The signals may enter the SDH hierarchy at any of these levels and be nested into the Synchronous Transport Module level 1 frame (STM-1). This is the basic unit of bandwidth demand that the core transport network would deal with, approximately 155 Mbits/s.

As the application interfaces evolve, for example in storage area networking using 1Gbit/s and 2 Gbit/s Fibre Channel, and Gigabit Ethernet interfaces become more widely adopted there is more potential for the transport hierarchies to become compressed such that the service rates are a significant proportion of the transport rate, for example 1 Gbit/s or 2.5 Gbit/s in 10 Gbit/s.

2. The model framework

The network model framework is most efficiently described referring to figure 1 This outlines the network aggregation hierarchy and demonstrates how the traffic demands penetrate the network depending upon the "distance" between the source and sink nodes. The geographic aggregation may be varied independently of the line aggregation to allow "statistical multiplexing" of the traffic loading.

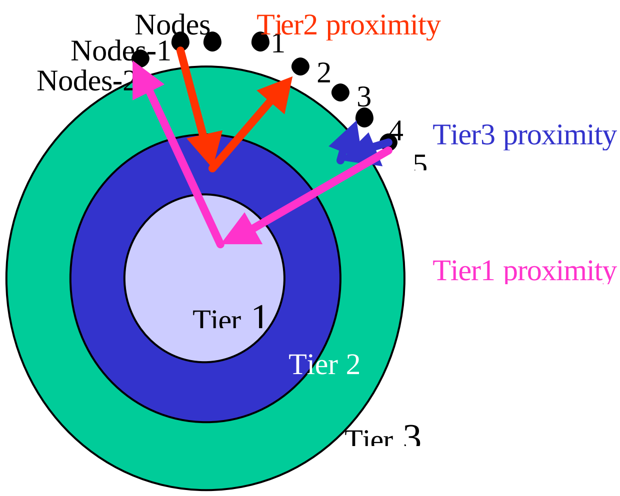

Figure 1.

Network model schematic

The ratio at which traffic demands are created and destroyed defines the amount of traffic “churn” in the network. Both the traffic sources and sinks are randomly generated, creating a random mesh loading on the network.

3. The Network Hypothesis under test.

We wish to test the hypothesis that under conditions of network traffic growth and churn, where the network has more than one level of traffic aggregation, the opportunity to reuse bandwidth (wavelengths) within the core of the network will be minimal.

Thus the growth and churn take place at tier 3 in the network, aggregation takes place at tiers 2 and 1 and the effects of interest are taking place at tier 1.

This impacts the level of flexibility that network equipment may require, and ultimately both the hardware costs (CAPital EXpenditure) and running costs (OPerational EXpenditure) of the network.

4. Model Conditions

Two aggregation hierarchies are of immediate interest the x4 of STM/SONET and the x10 which is used in Ethernet. The optical line rates of interest are 1GE (1.25 Gbit/s) STM-16 (2.5 Gbit/s) and STM-64/10GE (10 Gbit/s). It should be noted that traffic grooming is assumed as transitions between the network tiers will invoke a different line system, to accommodate the increasing transmission distance involved.

The traffic churn / growth ratio range of interest is 1%-25%

5. Results

Figure 2 illustrates the behaviour of the model when examining the potential for wavelength deletions in the core of the network (tier 1). This result is generated from 100k tier 3 (STM-1) demands aggregated through STM-16 to the STM-64 figure shown. The random nature of the demand pattern can be seen together with the empty leading diagonal, as traffic is not routed to and dropped from the same tier 1 node. The scale of this network is representative of a U.S. or pan European network with approximately 25 major nodes supported by a 40-wavelength line system.

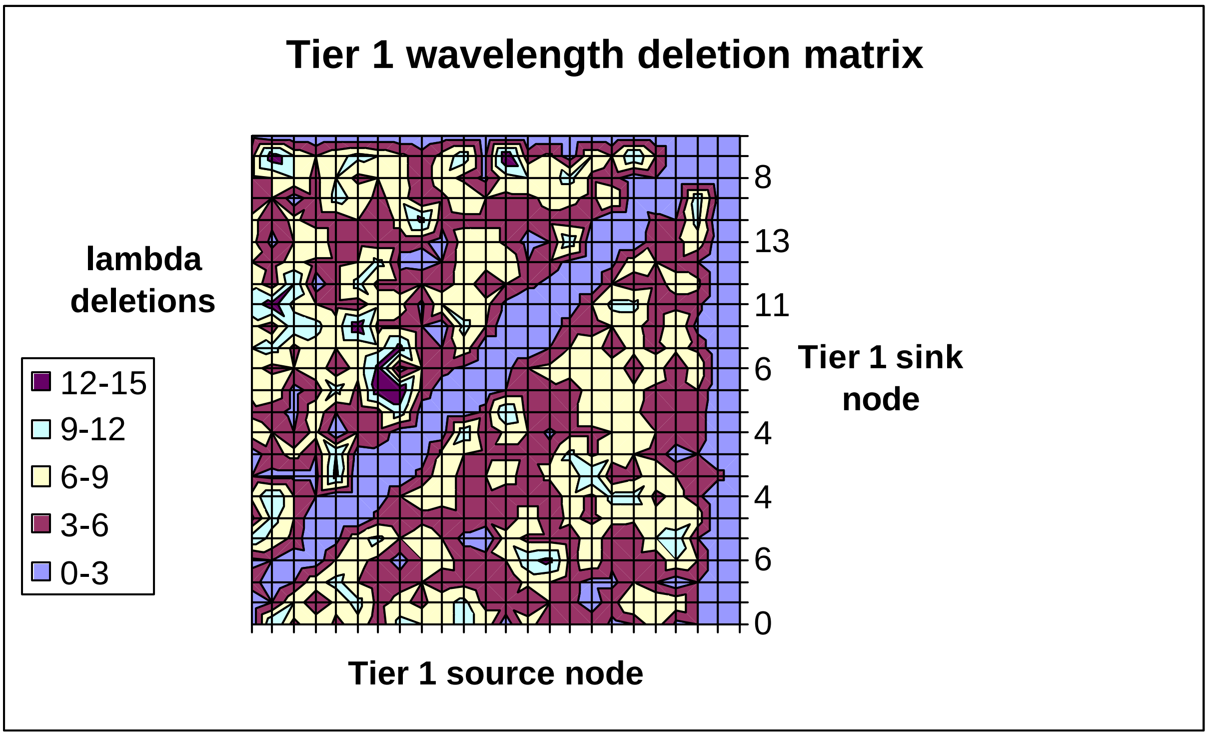

Figure 2.

Tier 1 wavelength deletion matrix

In order to quantify the scale of wavelength deletions with changes in the traffic “churn” (delete/add ratio) a series of simulations was conducted noting the maximum, mean and minimum number of wavelength additions and deletions. Figure 3 summarises the results. The minima in all cases were zero and have been omitted form the figure.

It can be seen that the maximum number of wavelength deletions increases with traffic churn under the conditions that a new wavelength is added when the previous wavelength is 75% full and deleted when the penultimate wavelength is 50% full. See the discussion section for the rationale of these numbers.

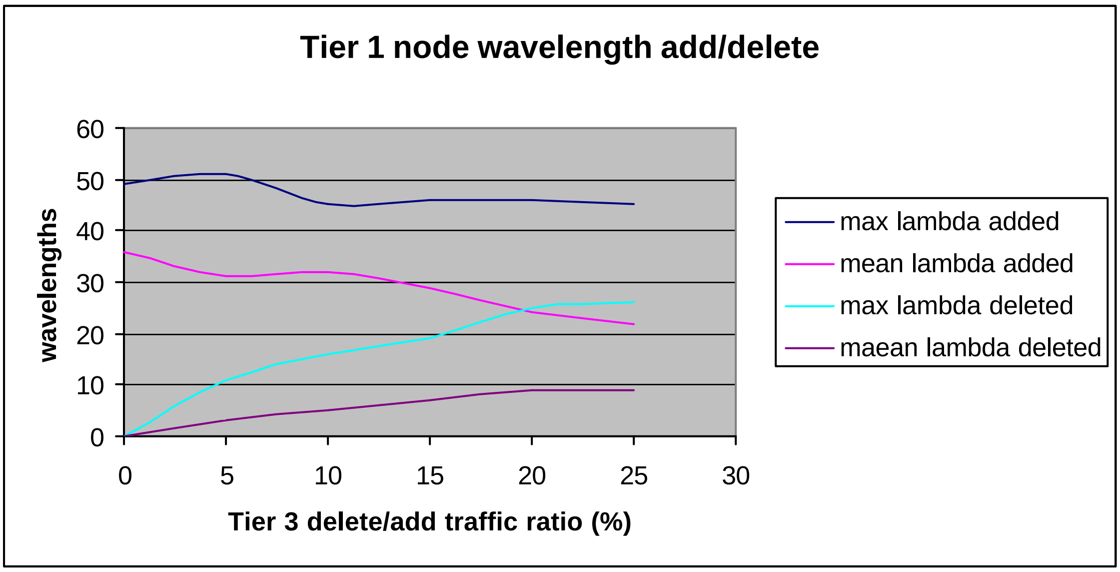

Figure 3

Wavelength evolution wit h churn

As a proportion of the maximum number of wavelengths being deployed in the network the maximum number of deleted wavelengths is given in figure 4.

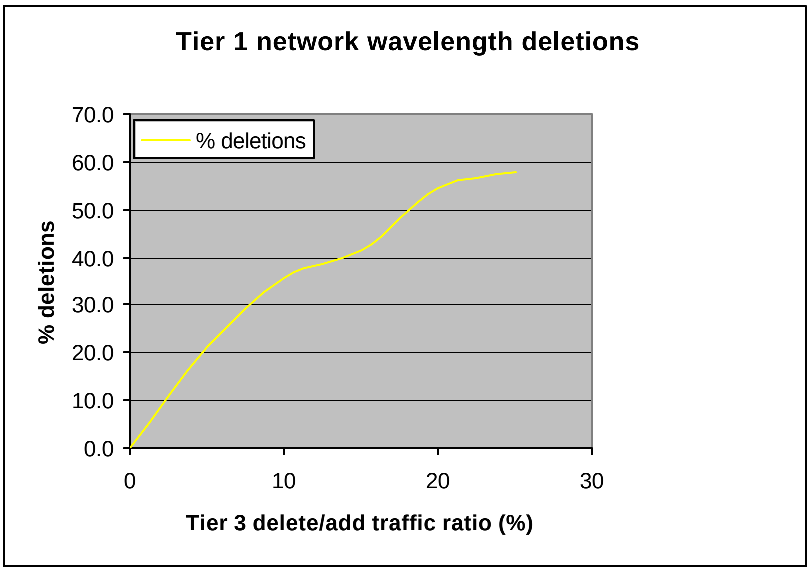

Figure 4

Maximum wavelength deletions as a percentage of maximum added traffic

6. Discussion

The pragmatic approach to ensuring that the capacity margin in any route is maintained is to have high and low water-marks that trigger maintenance events, adding or moving capacity at a node. High water-mark figures vary between 50%-90% depending upon the time taken to deliver the capacity. Low water-marks are less certain, this study has used 50% to supply some hysteresis to the process of wavelength add/delete.

When this hysteresis was increased to a full wavelength then the traffic deletion matrices, as shown in figure 1, emptied for churn ratios up to 30%. This behaviour will need to be checked in future investigations using other, less uniform, traffic generation profiles, as it may be due to the “spreading” of traffic with a random generator.

7. Conclusions.

The effects of traffic churn in the presence of overall traffic growth has been investigated under the conditions of random traffic generation. Taking a nodal view then wavelength deletions are comparatively rare events with respect to additions, however network wide these sum to a significant number of events. Thus the hypothesis has not been validated.

However, the operational expenditure in evolving the network to accommodate the changing traffic will be minimised by delaying the potential redeployment of wavelengths at any particular node until traffic growth at that node either reuses the wavelength or triggers the need for further wavelength deployment. The study indicates that delaying removal of a wavelength until the penultimate wavelength is empty (i.e. one wavelength of bandwidth) hysteresis) will decrease the wavelength deletion events to near zero. Under these operating conditions the hypothesis is supported.

Acknowledgments.

I would wish to thank Nortel Networks, UCL and EPSRC for supporting this work.

References.

[1] Internet Traffic Soars, But Revenues Glide,Muayyad Al-Chalabi, Director, Executive Strategic Program, RHK May 6, 2002, http://networkwatch.rhk.com/ShowAnalysis.asp?analysisId=108