NTuA3.pdf

Transponder Pool Sizing in Highly Dynamic Translucent WDM Optical Networks

Ronald A. Skoog, Brian J Wilson
Applied Research, Telcordia Technologies, Inc., 331 Newman Springs Road, Red Bank, New Jersey, 07701
{rskoog, brian2}@research.telcordia.com

Abstract: We have developed a methodology for sizing optical network node transponder pools to meet network call blocking requirements. The methodology is practical and adaptable for use in virtually any carrier's optical network capacity planning process.

OCIS codes: (060.4256) Networks, network optimization; (060.4253) Networks, circuit-switched

1. Introduction

The DARPA CORONET program [1] has laid out a vision for a highly dynamic, Wavelength Division Multiplexing (WDM) optical networking environment for providing on-demand wavelength services. A major challenge is to provide wavelength service connections using as few conversions between the electrical and optical domains as possible. The transponder (TxRx) is the elemental device that converts signals between the electrical and optical domains. At the edge of a long-haul WDM optical network, Add/Drop ports are used to convert client signals between the electrical domain and the WDM optical domain. For each client signal this Add/Drop function is done with a single transponder. The transmitter (Tx) maps an incoming electrical signal to a desired outgoing WDM wavelength (λ); and the receiver (Rx) maps the WDM λ optical signal to an electrical signal. In the CORONET environment the transponders are tunable, so each transponder can be used to map signals between the electrical domain and any of the allowed WDM wavelengths (in CORONET there are 100 wavelengths per fiber).

In addition to the Add/Drop functionality, there are two other uses for the transponder in an optical network. One is to do wavelength conversion, and the other is for regeneration (correcting for loss and impairment accumulation). Both regeneration and wavelength conversion are accomplished with two back-to-back transponders (electrical-to-electrical connection for both directions: Rx $\rightarrow$ Tx, Tx $\leftarrow$ Rx).

Transponders are by far the largest cost component in a WDM Optical Network, and thus it is essential to minimize their use as much as possible. Networks that are designed to maximize all-optical switching and minimize transponders are called translucent networks. The most efficient way to provide transponders in a translucent network is to have shared pools of transponders at strategically selected nodes in the network [2]. By "shared transponder pools" we mean the optical switches are configured so that every deployed transponder in the switch can be used for any Add/Drop, wavelength conversion or regeneration functionality (e.g., see Fig. 1 in [2]). Thus, in a translucent network, a key network design problem is determining at which switching nodes to locate transponder pools, and sizing the transponder pools to meet call blocking requirements.

This paper is concerned with the problem of sizing transponder pools to meet call blocking requirements, once it has been determined where to locate the transponder pools. In the CORONET network, there are 100 nodes globally, and 40 of those nodes support Wavelength Service (WS) Add/Drop. Those 40 WS Add/Drop nodes by definition must have TxRx pools for Add/Drop, and therefore those TxRx pools are also used and sized for providing wavelength conversion and regeneration functionality. In a separate study (not discussed here), 13 other nodes were identified to support TxRx pools for just wavelength conversion and regeneration functionality.

There are a number of previous studies that address the problem of locating and sizing transponder pools for translucent networks (e.g., see [2–4] and their references). However, those contributions develop methods that are only applicable for small networks and placing small numbers of transponders (e.g., less than 10 at any node). The CORONET program defines four network scenarios that have nominal wavelength service loads ranging from 5 to 25 Tb/s, fibers carrying 100 wavelengths, and wavelength bandwidths of 40 and 100 Gb/s. The previous methods for placing and sizing transponder pools are based on methods that cannot handle these carrier network scales requiring on the order of 40-50 transponders per node and 2-3 thousand transponders network wide.

2. Transponder Pool Sizing Methodology

The input to the transponder pool sizing process is the network topology (nodes, links, and fiber-pairs {1 or 2} on each link), the wavelength service traffic intensity matrix between the 40 wavelength service nodes, wavelength service call properties (proportion of calls requiring 1, 2, 4 or 8 λs; call holding time distributions), and the nodes in addition to the 40 wavelength service nodes that are used to support transponder pools (identified by other studies). Wavelength service call simulations are run with "unlimited" (e.g., 1000) TxRx in each transponder pool. For each transponder pool, samples are periodically taken of the number of TxRx in use. The samples are separated by an interval of around 30 minutes simulated time to get reasonably independent samples.

The result of the above simulation and transponder pool sampling process is to obtain histograms for each transponder pool. The basic question we explored is whether there is a predictable probability distribution for each transponder pool that can be used to size the pool for a desired blocking probability. It turns out that each TxRx pool does indeed have a predictable distribution. We found that all of the distributions are members of the Chi-Squared (χ^2) family. The Chi-Squared distribution family is a one parameter family having distribution function illustrated in Figure 1. The parameter *k* is a positive integer, called the degrees of freedom. Different transponder pools can, and do, have different χ^2 degrees of freedom.

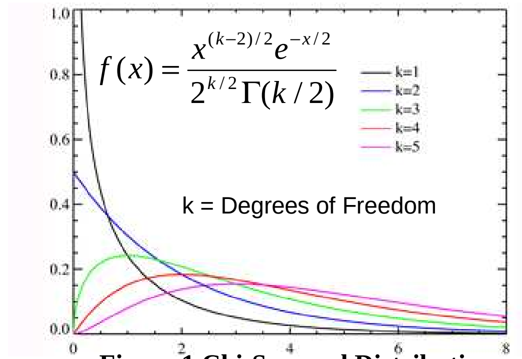

Figure 1 Chi-Squared Distribution

In order for this to be a predictable and reliable process, it is necessary that for a given network topology and traffic intensity matrix, the χ^2 degrees of freedom for each transponder pool should be the same (or within ±1) for different simulation seeds. In the results discussed below, we found this to be the case. Therefore, having known, stable χ^2 distributions for the histogram data, we sized each TxRx pool so the tail area of its χ^2 distribution to the right of the pool size would equal a desired TxRx related blocking probability. As discussed below, the predicted blocking probabilities closely matched simulation values. It is important to note, however, that the tail area of a χ^2 distribution beyond some value (or beyond some number of standard deviations) varies measurably with its χ^2 degrees of freedom, so it is critically important to accurately identify the χ^2 degrees of freedom for each TxRx pool.

3. Statistical Analysis of TxRx Pool Histogram Data

Figure 2 illustrates the type of statistical analysis that was done to validate that a TxRx pool histogram has a χ^2 distribution and determine its degrees of freedom. We used the NIST Dataplot statistical software [5] for the analysis. The left box in Figure 2 shows the histogram for Atlanta, which supports WS add/drop. The middle box shows a Probability Plot Correlation Coefficient (PPCC) Plot assuming the distribution is Chi-Squared. PPCC plots show for different values of a distribution family's parameter value how well that distribution fits (correlates with) the assumed distribution. This plot is used to find the best candidate (maximum correlation) for the χ^2 degrees of freedom. Then, to evaluate how well that distribution does fit, we use the Probability Plot, shown in the right box, assuming the distribution is χ^2 with the previously determined degrees of freedom. The Probability Plot essentially plots the cumulative distribution of the histogram under test (vertical axis) against the cumulative distribution of the assumed distribution (horizontal axis). If the plot is linear, the histogram under test fits the assumed distribution. Deviations from linearity indicate deviations from the assumed distribution. The slope gives the scaling and the intercept gives the translation that must be applied to the assumed standard distribution for it to correspond to the data.

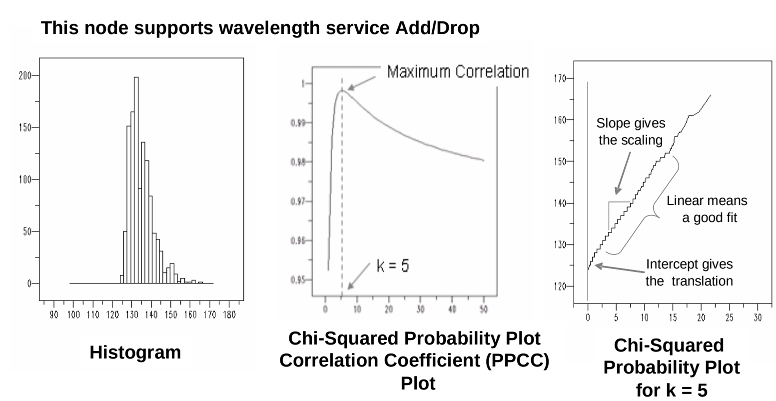

Figure 2 Dataplot Statistical Analysis of TxRx Pool Histogram

4. Simulation Results

In order to test the methodology described above for sizing TxRx pools, we used the simulation capabilities developed by the CORONET-funded PARAGON (Protocols, and Architectures for Resilient, Agile, Global, Optical Networks) project. The CORONET optical network consists of 100 nodes and 136 links. CORONET defines four network bandwidth scenarios (with 20, 40, 50 and 100 Tb/s of total network traffic, respectively). In this study, we considered Scenarios 1 and 2 (20 & 40 Gbps). In Scenario 1 there is one fiber pair per link, and 11 links were augmented with an additional fiber pair. In Scenario 2, there are two fiber pairs per link, and 2 links were augmented with an additional fiber pair.

For each scenario we ran 5 simulations for TxRx pool sizing and 5 for blocking, using a different seed for each simulation, and each simulation covered 2 weeks of simulated time. Two wavelength service classes dominate the call blocking. One has a maximum holding time of 1 minute and the other has an average holding time of about 3 hours. Thus, 2 weeks of simulated time gives a significant turnover of calls. Wavelength service calls can be for 1, 2, 4 or 8 λs. The nominal amount of bandwidth used by wavelength services is 25% of the total network bandwidth (5 Tb/s for Scenario 1, 10 Tb/s for Scenario 2). Forty percent of that bandwidth is consumed by single λ calls, and 20% each is consumed by 2, 4, and 8 λ calls. From that data it can be determined that there are on average 1.74 λs per call.

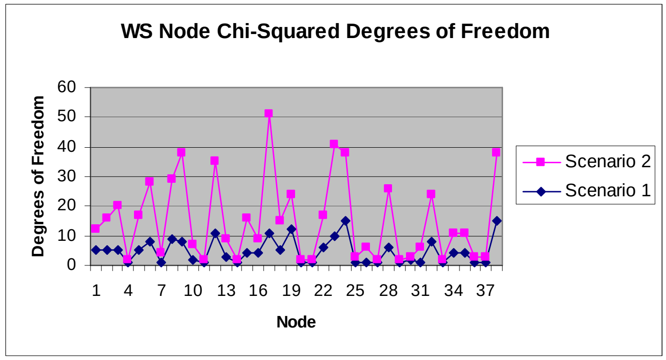

Figure 3 Node χ^2 Degrees of Freedom for Scenarios 1 and 2

To size the TxRx pools for each scenario, we combined the samples from 5 simulations (seeds) for each node supporting a TxRx pool, and analyzed those sample histograms as described above. We showed that all those nodes had a χ^2 distribution, with a significant range in the degrees of freedom. Figure 3 shows the degrees of freedom obtained for the nodes supporting WS Add/Drop. There is a significant change for most of the nodes when going from Scenario 1 to Scenario 2, showing that traffic intensity and topology changes (number of fiber pairs on links) can significantly change the χ^2 degrees of freedom at a node, but not the fact that the distribution is χ^2. For Scenario 1, all the nodes with TxRx pools not supporting WS Add/Drop have χ^2 degrees of freedom equal to 1. For Scenario 2 all but 3 nodes have degrees of freedom equal to 1; for the other three it is 2.

We sized the TxRx pools using a 10^{-4} tail area on the χ^2 distributions, an approximation for achieving 10^{-4} TxRx blocking probability. As discussed above, there are on average 1.74 λs per call, and from simulations we determined that for both Scenarios 1 and 2 there are an average of 3.2 TxRx used per λ connection (2 for Add/Drop and 1.2 for regeneration or λ conversion). Thus, there are on average 5.57 (3.2x1.74) TxRx per call, and a 10^{-4} blocking probability for transponders would imply a 5.57×10^{-4} transponder- related call blocking probability. To determine the TxRx related call blocking from simulation, we ran simulations to get call blocking when there are no TxRx limits (1000/node) and simulations with the TxRx pool sizes described above. The difference in those blocking results gives the TxRx related call blocking. For Scenario 1 this resulted in a 5.74×10^{-4} TxRx blocking ratio, and for Scenario 2 it was 3.23×10^{-4}. These are relatively close to the expected value of 5.57×10^{-4}!

5. Conclusions

We have developed a methodology for sizing TxRx pools in translucent optical networks that is applicable to realistic carrier scale networks. The methodology uses well defined simulations using data that a carrier would have available as part of their capacity planning process. The statistical analyses required could be automated, but we have not explored that aspect. An intriguing question is why all the TxRx pool histograms fit a χ^2 distribution.

Acknowledgement
The support of the DARPA CORONET Program, under DARPA Contract N00173-08-C-2011, is gratefully acknowledged.

References

[1] DARPA, "www.darpa.mil/STO/Solicitations/CORONET/index.htm", 2006.
[2] X. Yang, B. Ramamurthy, "Dynamic routing in translucent WDM optical networks," 2005, IEEE JLT, **23** (3), pp.955-971
[3] X. Chu, B. Li, I. Chlamtac, "Wavelength converter placement under different RWA algorithms in wavelength-routed all-optical networks," IEEE Trans. Comm., 2003, **51** (4), pp. 607-617.
[4] X. Hei, et al, "Wavelength converter placement in least-load-routing-based optical networks using genetic algorithms," 2004 JON, **3**(5), pp 363-378
[5] Heckert, N. A. (2001). Dataplot website at the National Institute of Standards and Technology: http://www.itl.nist.gov/div898/software/dataplot.html/, June 2001