TWO-TIER NETWORK ECONOMICS

Christopher Olszewski, Eric Bouillet, Jean-François Labourdette,
Sid Chaudhuri, Don Smith, and Yiwen Tan
Tellium, Inc., 2 Crescent Place, Oceanport NJ 07757
colszewski@tellium.com (corresponding author)

Introduction

An efficient and well-operated network is designed so that its parts function well together. This coordinated efficiency arises from having each element of a network perform those tasks for which it is best designed. Similar to the division of labor in a human society, the division of labor amongst elements in a communications network allows them to maximize the network's output (e.g., bandwidth connectivity, restoration speed) at the minimum cost.

A multi-tier (or layered) network provides economic efficiency in such a way: each tier performs tasks that it can handle efficiently and economically; tasks that one tier cannot efficiently or economically perform are left to another that can perform those functions effectively. In a small network, it may be economically justifiable to have fewer elements performing duties less efficiently, because the small size of the network does not unduly penalize the inefficiencies. But, in a large network, the efficiencies of specialized tiers will likely outweigh the cost of the extra network elements.

This paper investigates the economics of a two-tier division in a modern transport network: a core transport tier, and an edge grooming/traffic aggregation tier. The cost of this two-tier transport network is then compared to the cost of a network made of a single tier. Comparisons are limited to the cost of the transport capacity needed for the network and to the restoration performance of the network. Other expected benefits of a two-tier network, e.g., operational savings and improved reliability, will not be directly addressed.

Two-Tier Network Motivation

Based on historical trends, current and projected growth of data services, and improved restoration speed, the core switching rate of 2.5Gbps is the next logical step. This paper investigates the economic viability of such a 2.5Gbps tier in a two-tier network. We compare a one-tier network operated at the STS-1 level with a two-tier network operated at the STS-48 level in the one tier, and at the STS-1 level in the other tier.

Historical Trends

As Digital Cross-Connect Systems (DCSs) have been introduced into digital core transport networks over the past 25 years, the rate of the core transmission speed has traditionally been about 20 to 40 times the rate of the core "switching" (cross-connect) rate.[1] DS0 (64 kb/s) signals were switched when core transport systems were on the order of DS1 signals (1.5 Mb/s); similarly, DS1 signals were switched within DS3 (~45 Mb/s) signals. Most recently, DS3 signals were switched with DCSs when the core transmission speeds were on the order of 1.5 Gbps to 3 Gbps (~STS-48). Now, when the core transmission speeds are on the order of multiple 10Gbps wavelengths, by historical analogy the rate of core transport switching should be on the order of an STS-48 (2.5Gbps).

Current Networks

Client devices are now being produced and installed with interfaces at the rates of 2.5 Gbps and 10 Gbps.[2] As these devices are deployed more frequently, it makes sense that these customer interface rates also become the rates for the core switching tier of the transport network. For these rates, maintaining the capability to switch traffic with more granularity is inefficient.

Restoration Performance

Given similar restoration architectures, the speed of restoration (e.g., number of STS-48 services restored per 100 ms) will be greater within the STS-48 tier of a two-tier network than within the STS-1 tier of a single-tier network.[3] The greater performance in the STS-48 tier is due to that tier's exclusive use of STS-48 and above switching. Using a shared-mesh restoration scheme,[4,5] the failure of a single STS-48 service will require one switch to take place at each node along a shared protection route. In this network, restoration of services after a reasonably-sized fiber cut will likely take place within a timescale on the order of 100 ms. Transport restoration times of 200 ms or less should have minimal impact on carried services.[6]

In contrast, STS-1 switches in a single-tier network must make connections on an STS-1 by STS-1 basis. Thus, for a single STS-48 connection requiring restoration, the STS-1 tier will make about 50 STS-1 connections per switch along a shared backup path. This scaling argument suggests that restoration speeds will be about 50 times slower for an STS-1-based network than for an STS-48-based network. Restoration times are not strictly linear with the number of entities being restored, and so this estimate may be optimistic. Hence, the shared-mesh restoration performance of an STS-1-based tier will be on the order of seconds, at best.[3] Failures of this duration can disrupt existing services.[6]

Even in a dedicated backup path scenario (e.g. 1+1), with connections already established along the backup route, the restoration (or protection) speed of the STS-48 tier will be greater than that of the STS-1 tier because of the more granular nature of the STS-1 switch.

Because of its faster performance for equivalent network capacity deployed, the STS-48 tier thus has an advantage in this respect over the single STS-1 tier.

Specification of the Problem

The metric for cost comparison between the single-tier and two-tier networks will be capital expenses (CapEx). Essentially, this analysis seeks to determine the cost of building the one-tier and two-tier networks. Other areas of economic comparison (operational expense (OpEx), improved revenue, new services, network management complexity, etc.) will not be addressed.

To fairly compare these two network architectures, their relative restoration performance should be considered. The costs of two networks providing the same restoration performance form a valid basis for comparison. For restoration times on the order of 100 ms, we can compare the two-tier network (using shared mesh restoration in the STS-48 tier, and 1+1 restoration in the STS-1 tier), and the single-tier network (using just 1+1 protection). The two-tier network is expected to require less bandwidth than the one-tier network.

A single-tier STS-1 based network using shared mesh restoration would generally (but not absolutely) be expected to use fewer network resources than the two-tier network treated here. However, as previously stated, this economic saving is achieved at the expense of potentially disrupting current services during a failure. This disruption is not considered acceptable and, therefore, the single-tier shared-mesh restoration architecture is not considered further in this paper.

Network Architecture

Two model architectures were used in this analysis. Figure 1 shows a sketch of the one-tier network architecture. Ubiquitous STS-1 switches in different offices are connected through physical links. DWDM systems carry the physical connections from one node to another. For the purpose of this investigation, we assume that the STS-1 switches can restore using an end-to-end 1+1 mechanism.

In the two-tier network architecture, shown in Figure 2, the node architecture is not uniform. All nodes still have an STS-1 level switch, which is essential for handling sub-STS-48 traffic. However, some nodes, either carrying a large amount of add/drop traffic or occupying strategic switching locations, are called hub nodes and also have an STS-48 level switch. In this architecture, nodes without an STS-48 level switch are called non-nub nodes. For the purposes of this investigation, we assume that the STS-48 switches are capable of shared mesh restoration, while STS-1 switches are required to use an end-to-end 1+1 restoration mechanism to have comparable restoration speed.

We consider the architecture shown in Figure 2 to be a "mixed two-tier architecture", because not all nodes in the architecture are of the same kind: some nodes are non-hub nodes, while other nodes are hub nodes. Nodes are connected not only through STS-48 switches, but also through neighboring STS-1 switches.

Connections between the STS-1 and STS-48 switch in a hub node are protected with a 1:N LAPS protection switch arrangement. A connection along its path is protected within each tier as in a separate restoration domain, and could be protected by multiple mechanisms: a shared mesh scheme in the STS-48 tier, a 1+1 scheme in the STS-1 tier, and a LAPS scheme between tiers within a single node. Because the STS-48 tier domain and the STS-1 tier domain meet at a common transition point between the two domains, this overall protection scheme is sometimes referred to as a "pinched" model of protection.

In both single and two-tier networks, we assume that the DWDM is capable of only 2.5Gbps channels. Changing this assumption does not affect the general results of this study.

In the two-tier network, STS-48 switches are connected to their nearest neighboring STS-48 switches. Intervening STS-1 switches are not included in these connections because they are not able to participate in the same restoration protocols as the STS-48 switches. These "express links" connecting STS-48

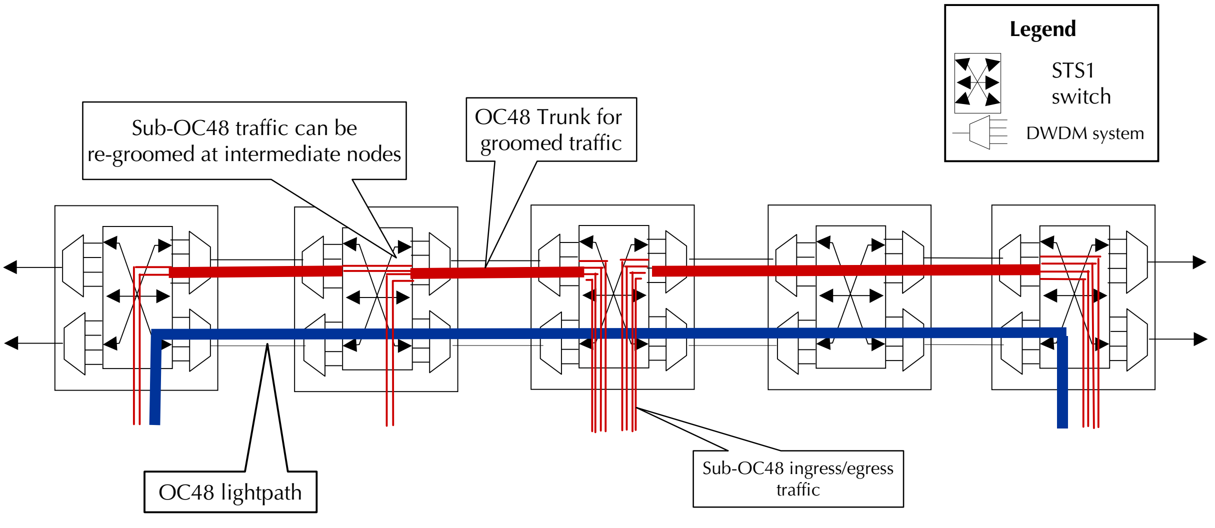

Figure 1 One-tier network architecture

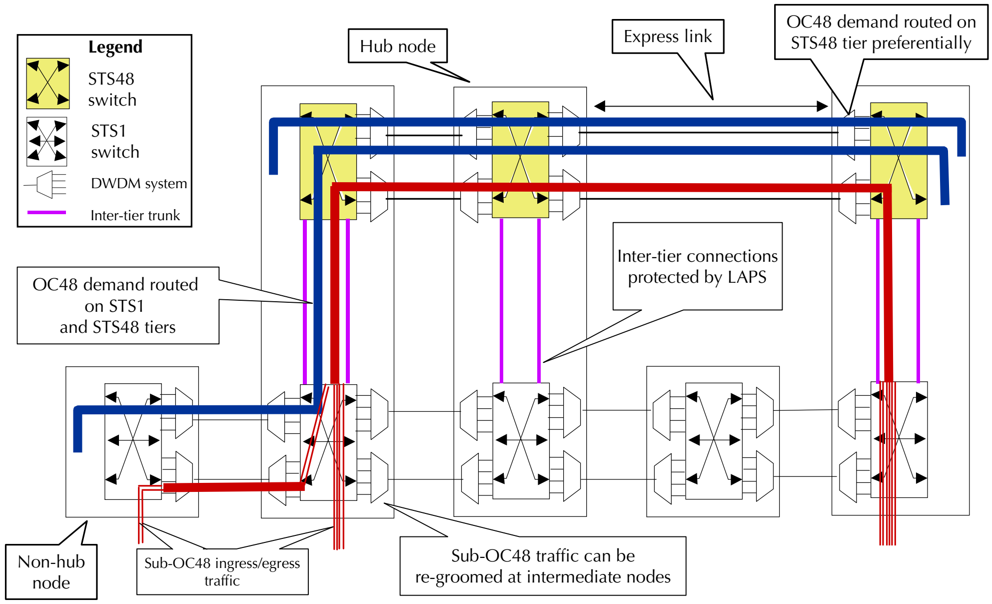

Figure 2 Two-tier network architecture

switches arise very naturally in the two-tier network (see Figure 2). In the one-tier network, such links could bypass several intermediate switches (and switch ports) with possibly considerable cost savings.[7] To compensate for this discrepancy, equivalent express links were added in the one-tier networks. In this architecture, both bypassed and terminating switches use the same restoration protocols.

Restoration Mechanism

In the single-tier network, the restoration mechanism is a dedicated 1+1 primary / backup pair of paths. This mechanism should be able to provide protection switching within 100 ms.

In the dual-tier network, the restoration mechanisms in each tier are logically separated from restoration actions in the other tier.

- Within the STS-48 tier, a fast shared mesh restoration scheme is assumed to act. (Because restoration at the STS-48 tier affects many fewer objects, it is at least an order of magnitude faster than a shared mesh restoration mechanism at the STS-1 level for the equivalent restored bandwidth). This mechanism is assumed to take place within 100 ms or 200 ms.[3]
- Within the STS-1 tier, the same dedicated 1+1 primary / backup protection mechanism is assumed to operate as in the single STS-1 tier.
- Between the STS-1 and STS-48 tier, a linear APS (LAPS) 1:N scheme is assumed to operate. For this study, N is taken as 8.

In this way, restoration/protection schemes are logically separate and self-contained within a tier, or between tiers. The choice of the 1+1 mechanism for the STS-1 tier in the two-tier architecture is made to give similar performance to the shared-mesh restoration scheme in the STS-48 tier.

Network Model

To investigate these economic factors in network deployment, we developed a model network. As a typical network of interest to US carriers, this model is of a national US network. The topology of the network was taken from public sources describing U.S. national networks and represents an amalgamation of existing networks.[8,9] Large population centers were also included. Topologies in some

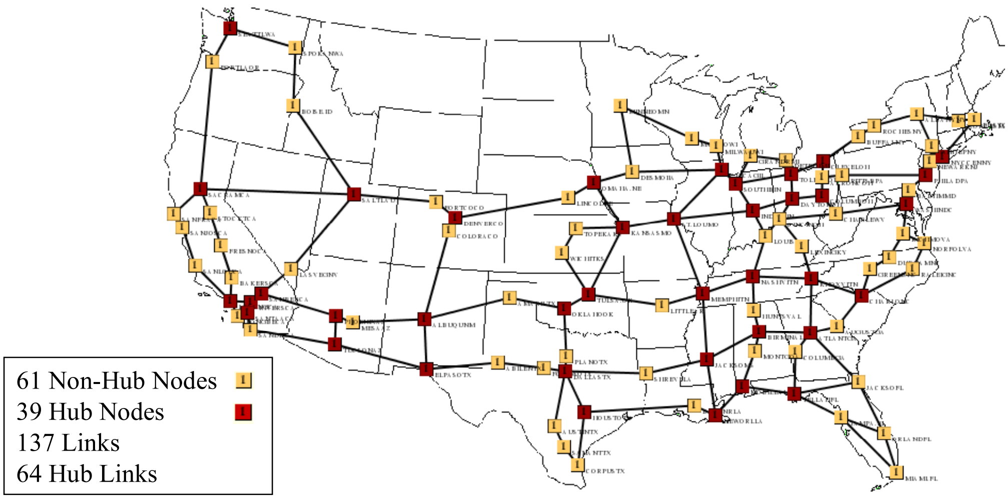

Figure 3 Model network topology

regions may differ between carriers, but for the most part there is a great deal of similarity in U.S. national networks. The base topology of this network is shown in Figure 3. There are 100 nodes and 137 physical links.

The traffic placed on this network consisted of demands at both the OC-48 rate and at sub-OC-48 rates. For the sake of simplicity, we assume that the sub-OC-48 rate traffic is all at the OC-3 rate. The balance between the amount of traffic at the OC-3 rate and at the OC-48 rate is set to be approximately 50% each, which is a reasonable assumption of many current and future networks.

We also characterize the nodes in the network as being large or small in terms of the traffic that originates at them. Large nodes have more traffic between them and other large nodes in the network. Small nodes typically have most of their traffic within a nearby geographical region. With this traffic pattern, we generated traffic on the model network. Table 1 shows a breakdown of traffic types and destinations. The overall normalization of the traffic is then a parameter that can be adjusted to raise or lower the overall traffic volume. For the purposes of this study, the overall normalization is fixed at about 3,600 OC-48 equivalents of traffic, which is consistent with that of a large national carrier. Varying the traffic volume by very wide margins does not change the main conclusions of this study.

Parameters

The parameters for the network model are consistent with market research reports and industry knowledge.[10] They are shown in Table 2.

Connections between the STS1 and STS48 tier are assumed to be at the OC-48 rate.

Table 1. Traffic concentrations for base case and diffused demands scenario.

Quantity	Base Case	Diffuse Demands
Large to large node traffic	76%	53%
Large to small node traffic	19%	41%
Small to small node traffic	5%	6%

Table 2. Model parameter values

Quantity	Value
2.5G port cost (both STS-1 and STS-48 switches)	$20K
2.5G transponder cost	$15K
WDM terminal	$250K
2.5G regenerator transponder cost	$18K
OA cost	$50K
Wavelengths on WDM system	80
OA spacing	50 miles
Regenerator spacing	350 miles

Because add/drop ports on the network are the same regardless of architecture, they are not considered in the cost model (except when calculating the overall number of ports and switches at a node). The cost of the switches is assumed to be fully amortized into the cost of the ports. The size of the STS-1 and STS-48 switches is initially taken to be infinite (i.e., only one switch required at any location). However, this assumption is adjusted to examine the effect of physical switch sizes (i.e., 512 and then 2048 OC-48 ports for the STS-48 switch, and 256 and then 512 OC-48 ports for the STS-1 switch).

Packing and Routing

In the one-tier network, because the switches are all capable of STS-1 switching, packing the traffic incurred virtually no penalty. All demands are routed along the shortest physical path, and then a 1+1 protection mechanism, or a shared-mesh mechanism, is applied to the resulting lightpaths.

In the two-tier network, STS-48 switches cannot groom traffic below the STS-48 rate, and so STS-1 switches are needed. The STS-1 switches groom traffic into STS-48 lightpaths, which are then routed with mesh protection across the network of STS-48 switches.

The traffic is groomed in the following simple way: end-to-end traffic that fills more than 70% of an OC-48 initiates the creation of an end-to-end OC-48 demand. That OC-48 demand carries the end-to-end sub-rate traffic.

End-to-end sub-rate traffic that does not fill an OC-48 is assumed to follow the least-cost physical route. Along this route, enough OC-48 demands are created on each link to carry all the traffic that is so routed.

The resulting set of OC-48 demands in the STS-48 tier are routed using shared mesh protection.

Base Case Results

The results of the base case models are shown in Figure 4, Figure 5, and Figure 6, for the number of network OC-48 ports, bandwidth miles, and costs, respectively. Figure 4 shows that, for a network offering ~100 ms restoration performance, the two-tier network uses about 13% fewer network ports than the one-tier network. Network ports include inter-tier (and inter-switch) ports, but exclude drop ports. The two-tier network also requires about 28% less bandwidth because of the shared mesh restoration scheme (Figure 5), and also costs nearly 20% less than the one-tier network (Figure 6).

The sizes of the switches in the various architectures are shown in Figure 7. Since the switches are assumed to have an unlimited number of ports, these distributions show the number of ports needed at each node. The distribution of switch sizes in the one-tier network shows a tail out to 2000 ports and

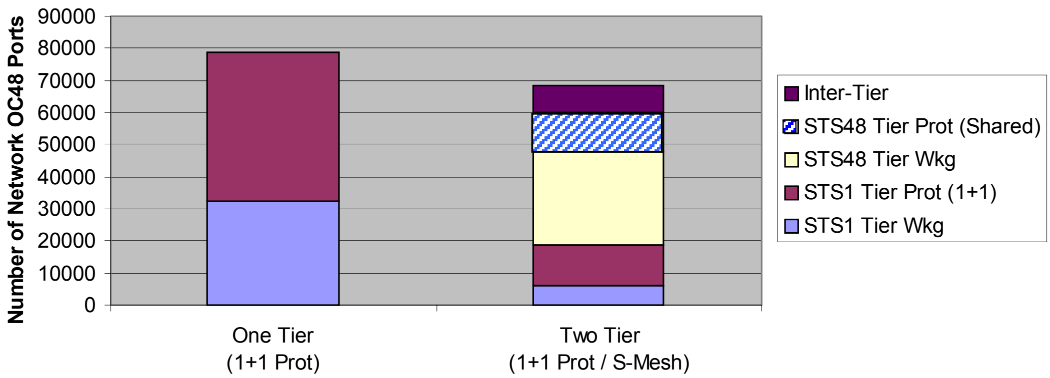

Figure 4 Switch ports for base case study

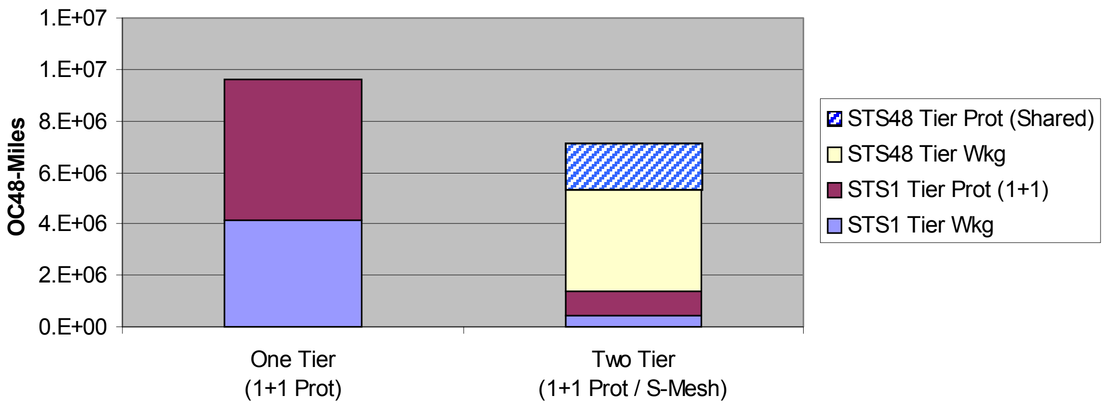

Figure 5 Bandwidth distance for base case study

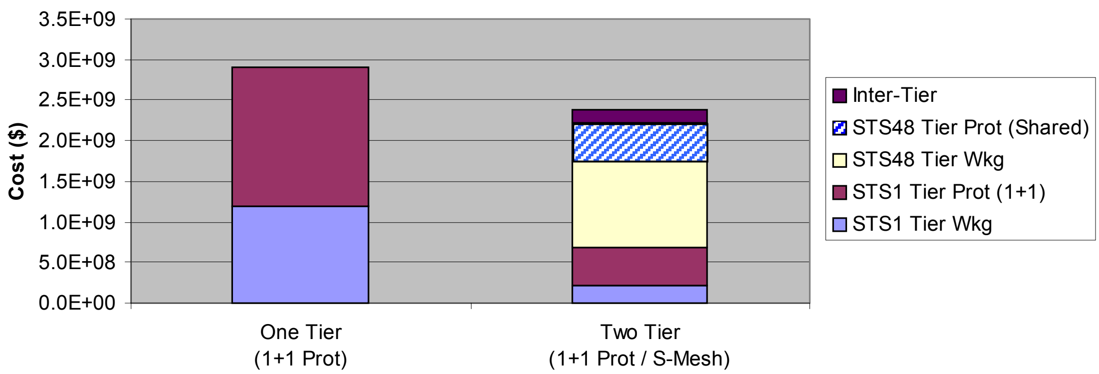

Figure 6 Network costs for base case study

beyond. For the two-tier network, the STS-1 switch size is effectively less than 500 ports, while the STS-48 switch extends up to about 1800 ports.

Sensitivity Studies

We conducted several sensitivity studies to investigate these results with changes to our assumptions. These studies included:

1. Using a finite switch size
2. Changing the traffic distribution (diffusing the traffic over more nodes)
3. Varying the relative costs of the OC-48 ports (alternately making them 20% less expensive in the STS-1 tier, and then in the STS-48 tier)
4. Removing the express links in the STS-1 tier.

An investigation changing the overall amount of traffic (from –75% to +25%) found no dramatic changes to the conclusions, and so is not discussed further.

Summaries of all the sensitivity studies are given in Table 3.

Finite Switch Size

The base case analysis assumes that both the STS-1 and the STS-48 switches have infinite capacity; or, that multiple like switches in a single location can be connected together without penalty. This section examines the effect of having a more realistic number of ports on the switch. The sizes of the switches in the base case are shown in Figure 7, which does not account for switch size limitations.

To incorporate the finite switch size and the subsequent need to interconnect switches within an office, a switch port penalty function is applied to locations that exceed the size of one switch. This function attempts to reflect the need to have cross-tie connections between multiple switches in an office, to accommodate discrepancies between forecasts and actual demands, and to improve sharing of restoration facilities.[11,12] Based on some simple analysis, this penalty function is defined as:

- For locations requiring one switch: no penalty
- For locations requiring between one and two switches: a 20% penalty based on the number of ports beyond one switch (maximum penalty: 12.5%)
- For locations requiring between two and three switches: a 33% penalty based on the number of ports beyond one switch (maximum penalty: 22%)
- For locations requiring four or more switches: a 50% penalty based on the number of ports beyond one switch (maximum penalty: 37.5% for four switches).

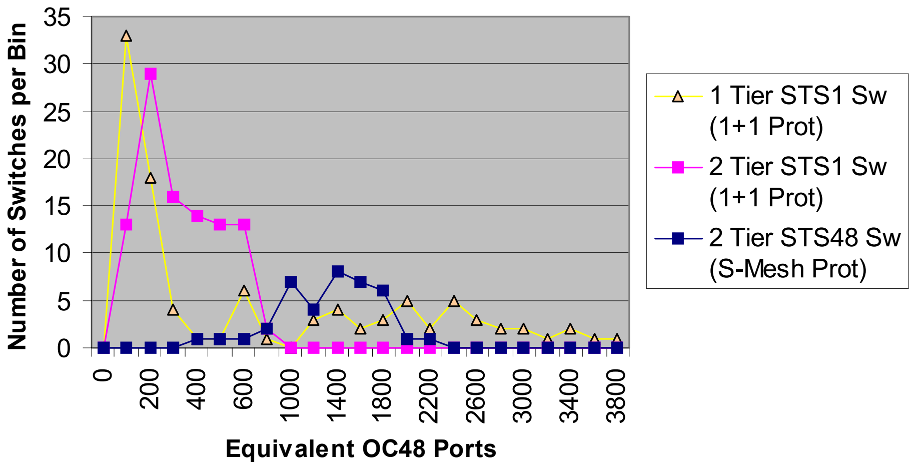

Figure 7 Switch port distribution for base case study

Table 3. Sensitivity study of network costs.

Sensitivity Study	1-Tier (1+1) ~100 ms Restoration	2-Tier (Shared / 1+1) ~100 ms Restoration
Base case	$2,912M	$2,387M
Finite switch size (STS-1: 256/STS-48: 512)	$3,615M	$2,762M
"Nominal case" Finite switch size (STS-1: 512/STS-48: 2048)	$3,503M	$2,396M
Diffuse demand set	$2,801M	$2,527M
Relative switch cost (less expensive STS-1 switch ports)	$2,518M	$2,273M
Relative switch cost (less expensive STS-48 switch ports)	$2,912M	$2,162M
Express links removed from one-tier network	$4,030M	$2,387M

The effect of this penalty function is to require even more ports for interconnections at those locations having more ports than are available on one switch. Table 3 shows the effect on network cost of currently available switch sizes (256 OC-48 ports on an STS-1 switch, and 512 OC-48 ports on an STS-48 switch), and for hypothetical switch sizes of future products (a 512 OC-48 port STS-1 switch, and a 2,048 OC-48 port STS-48 switch).

Table 3 shows that, at least for the hypothetical future switch sizes, there is virtually no penalty in the two-tier network. However, the penalty for the one-tier networks is substantial (~20%). The use of the term "nominal" case in this context is described later. For the currently available switch sizes, the costs of both networks increase, but the cost of the two-tier network increases proportionately less than the one-tier network.

This relatively smaller increase in the two-tier network is due to the separation of functions between the switches (grooming in one device, long-haul transmission in another), allowing each switch to be smaller than a single switch performing both functions in the one-tier network. This separation of functions enhances scalability and the capability of the network to grow more easily.

Diffuse Demand Set

This section discusses the relative concentration of the traffic. Traffic going exclusively between large nodes was redistributed to go between large nodes and small nodes, and to go between small nodes. Table 1 illustrates the way that the traffic demands were changed.

The effect of this change on network costs is shown also in Table 3. As can be seen, the two-tier network is still less expensive than the one-tier 1+1 network with comparable restoration performance, but the amount of savings has been nearly cut in half, from 18% to 10%. This result is consistent with intuition: if traffic is mainly concentrated between hub nodes (practicing shared mesh restoration), then the shared mesh restoration, with its efficient capacity use, will be responsible for most of the restoration. If traffic is not concentrated as much at the hub nodes, then additional resources (protected with the less efficient 1+1 scheme) are needed to either bring the traffic to a hub node, or connect the traffic directly.

Not surprisingly, the costs of the one-tier network does not change very much with the diffusion of the traffic to non-hub nodes. The restoration mechanism in this architecture is essentially indifferent with regard to the presence or absence of hub nodes, and should be relatively insensitive to traffic distribution and concentration.

Relative Switch Costs

This section looks at the effects of changing the relative costs between the STS-1 and the STS-48 switches. The cost of the OC-48 ports on the STS-1 switches was decreased by one quarter. The resultant network costs are also seen in Table 3. They show that even with this reduction in the STS-1 switch port costs, the two-tier network is still about 10% less expensive than the one-tier network with 1+1 protection.

The inverse cost structure was also applied: the cost of the OC-48 port on the STS-48 switch was made 25% less expensive than the port on the STS-1 switch. Table 3 shows the resulting savings increases to 26% of the two-tier design relative to the one-tier design. This cost reduction could effectively arise by integrating the DWDM transponder onto the switch interface card.

This sensitivity study shows that even under a wide swing in the relative cost of the components (±25%), the two-tier architecture with mesh restoration is still less expensive than the one-tier architecture with 1+1 dedicated protection.

Express Links

The results shown so far include express links in the one-tier solutions. These express links are a natural part of the two-tier solution, since in the two-tier architecture not every node has an STS-48 switch: in order to effect restoration, the switches must be connected together. However, these links are not really an inherent part of the STS-1 single-tier network: that is, although these links reduce the cost of the network, that network can still function without these links.[7] In this study, they have been included to help isolate the effects of different factors on the network's costs.

The results of not including these express links in the network design are also shown in Table 3. Without the express links in the one-tier network, the two-tier network shows a 41% savings in cost (up from 18%) over the one-tier 1+1 network.

Using express links in the one-tier network would require some additional operational complexity, in order to allocate capacity between express and local links going between switches of the same type. These links are an intrinsic part of the two-tier network.

Discussion and Summary

From the architectures examined here, we take the scenario with the 512-port STS-1 switch and 2,048-port STS-48 switch as the "nominal" case. In this scenario, realistic size limitations and their effects are imposed on the designs. The port costs of the two switches are equal to one another, and both architectures possess the same connectivity (express links) in their networks. The nominal case thus represents a fair comparison of the one-tier and the two-tier architectures. In this fair comparison, the two-tier network presents 32% savings over the one-tier architecture.

Figure 8 shows a chart summarizing the cost savings described in this paper of the two-tier network over the one-tier network with 1+1 restoration. It ranges from about 3% (combining all the elements most unfavorable to the two-tier network) up to 54% (combining all the elements most favorable for the two-tier network). The base case of this study is shown with 18% savings, while the nominal case has a 32%

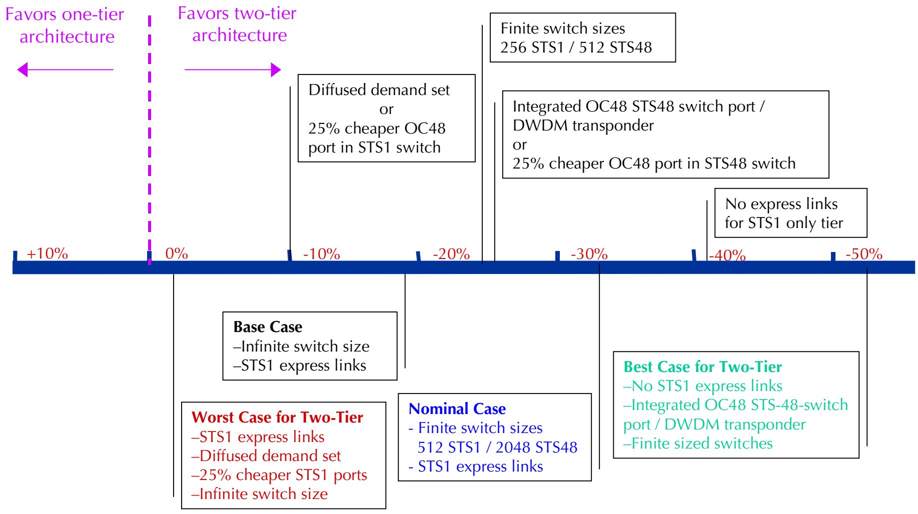

Figure 8 Summary of sensitivity studies

savings. Also shown in this diagram is the savings from integrating the OC-48 transceiver with the DWDM system. Given the range of the sensitivities for this comparison, this chart seems to indicate that the expected savings of a two-tier network are about 20% to 30% compared to the one-tier 1+1-protected network.

The sensitivity of the results to the finite switch size shows the good scaling properties of the two-tier network to network growth. The separation of functions to the STS-1 and STS-48 switches allows each tier to be optimized for its tasks, and keeps the switches smaller than in the one-tier cases. Accounting for the size of the switches as in the nominal case shows a savings of about 30% for the two-tier network.

The changes of the network costs due to the intensity of the traffic between the hub nodes indicate the need to choose hub nodes “well” in some sense. An interesting extension of this work would be to study the effects of different choices and numbers of hub nodes, perhaps as a function of traffic load.

Changes to the cost of the designs when the switch costs are varied indicate that, even within a broad range of relative price ratios, the benefits of a two-tier network still persist.

The most striking change in the sensitivity studies was the inclusion (or exclusion) of express links in the one-tier network. Even with these express links included in the one-tier designs, the savings benefits of the shared mesh in the two-tier network are considerable (18%). Without the express links in the one-tier network, the benefits of the two-tier architecture are unmistakable (~40%).

References

1. Stephen French, Jean-François Labourdette, and Krishna Bala, “Efficient Network Switching Hierarchy”, 2002 NFOEC proceedings.
2. Dell’Oro Group, Five Year Forecast WAN Infrastructure, January 2001.
3. Ahmet A. Akyamaç, Jean François Labourdette, et al., “Optical Mesh Network Modeling: Simulation and Analysis of Restoration Performance”, 2002 NFOEC proceedings.

4. G. Ellinas, E. Bouillet, et al., “Routing and Restoration Architectures in Mesh Optical Networks”, to appear in *Optical Networks Magazine*.
5. J. Labourdette, E. Bouillet, et al., “Routing Strategies for Capacity-Efficient and Fast-Restorable Mesh Optical Networks”, *Photonic Network Communications* journal, vol. 4, issue 3/4, 2002, special issue on *Routing, Protection and Restoration Strategies and Algorithms for WDM Optical Networks*.
6. ANSI T1.TR.68-2001, “Technical Report on Enhanced Network Survivability Performance”, Technical Report No. 68, February 2001.
7. Surendra Singh, Susan J. Chinburg, and Mark C. Wendel, “Efficient Network Design Using Mesh and Ultra Long Haul Configuration”, NFOEC 2001 proceedings.
8. Many references to maps of telecommunications networks can be found at “An Atlas of Cyberspace”, at http://www.geog.ucl.ac.uk/casa/martin/atlas/atlas.html.
9. National Communications System, “Management of Stressed Facility Networks, Level II Report Extension”, , Technical Information Bulletin 92-14, September 1992.
10. “Long-haul Transport & Optical Networks: North America, Market Forecast Supplement: Pricing Tables”, RHK, December 2001.
11. R.Ramamurthy, Jean François Labourdette, and Sid Chaudhuri, “Scaling Switching Capacity with Multiple Switches”, (to be published).
12. “Tutorial Corner”, Optical Networks Magazine, May/June 2002. Optical Network Magazine, tutorial corner, May/June 2002.