Benefits of Closer and Methods for Automatic Cooperation between Packet and Transport Networks

Dr. David E. McDysan
Verizon, 22001 Loudoun County Parkway, Ashburn, VA, 22407 USA
dave.mcdysan@verizon.com

Abstract: When Packet and Transport networks operate autonomously, performance degradations and inefficiencies can occur. Closer automatic cooperation between Packet and Transport networks results in improved performance and efficiency.
OCIS codes: (000.0000) General; (060.0060) Fiber optics and optical communications

1. Introduction

Traditionally, packet and transport networks have operated autonomously, with the packet network being a client of a transport server network, which can result in degraded performance and/or inefficiencies as described in Section 2. Closer automatic cooperation between Packet and Transport networks can result in a number of benefits, such as improved performance and efficiency as described in section 3. Generic requirements, potential methods and an overview of related standards for achieving cooperation are summarized in Section 4.

2. Problems with Autonomous Packet and Transport Network Operation

Problems which occur when packet and transport networks operate autonomously fall into two general categories: performance degradations, such as increased latency, and inefficiencies, usually resulting in higher cost.

Latency is a performance parameter which can be impacted by uncoordinated packet and transport network functions. We summarize several examples where this has historically been a problem in packet networks [2].

When a significant fraction of the traffic was private lines and voice trunks in the 1990s, using Synchronous Optical NETwork (SONET) rings as the restoration method for all circuit and packet traffic made good economic sense. But, to achieve the economic benefit of shared restoration, ring size had to be relatively large in a network with significant geographic extent (e.g., the continental United States). Latency around the shorter direction of a ring was good for packet networks, but during a SONET restoration action, latency around the longer direction of a ring could be significantly larger, affecting end user experience, sometimes violating quality objectives and/or resulting in user complaints in an Internet Protocol (IP) and/or MultiProtocol Label Switching (MPLS) network.

In general, packet networks deliver better performance (i.e., delay variation, latency and loss) when they are not heavily loaded. However, since restoration capacity is dedicated to rings, it is unavailable to the packet networksand this case of autonomous operation of transport ring protection and packet forwarding forced operation of more heavily loaded links, degrading performance. A response to this problem was to provision IP/MPLS networks over unprotected circuits and statically set the metric and/or Traffic Engineering (TE)-metric proportional to latency. This resulted in traffic being directed over the least latency path, but has a higher cost increased when compared with a design described in the next section where packet and circuit networks cooperate.

The economics of router interfaces have historically been driven by the market volume for interfaces of a particular protocol type. The wide scale usage of Ethernet electronics and optics beginning in the early 2000s means that routers will use Ethernet based interfaces for many years to come. Historically, Time Division Multiplexing (TDM) interfaces on routers have been more expensive, and hence a significant economic inefficiency occurs if routers use TDM interfaces.

Several decades of packet switched network operational experience have shown that statistical multiplexing gain is best achieved when the backbone speeds are much larger than the speed of the largest user packet flow. Since the header determines the destination in packet switching, it can dynamically adapt to short term variations in traffic patterns. A basic inefficiency results when routers are connected via circuit trunks of a speed lower than the speed of a router port, because unused capacity on a circuit toward a particular destination cannot be used by packets heading toward another destination. One way to quantify this inefficiency is to compute the maximum utilization as a function of circuit trunk speed using a statistical multiplexing model [1] as shown in Figure 1. The parameters chosen for this example are typical of a Fiber To The x (FTTx) or cable Internet access network with a user peak rate of 100 Mbps, an average user rate of 10 Mbps and a probability of queuing of 10^{-4}. As can be seen from this example, when the circuit trunk speed is only 10 times the user peak rate (i.e., 1 Gbps), maximum utilization is only 40%, while when the circuit trunk speed is 100 times the user peak rate, the maximum utilization increases to 70%.

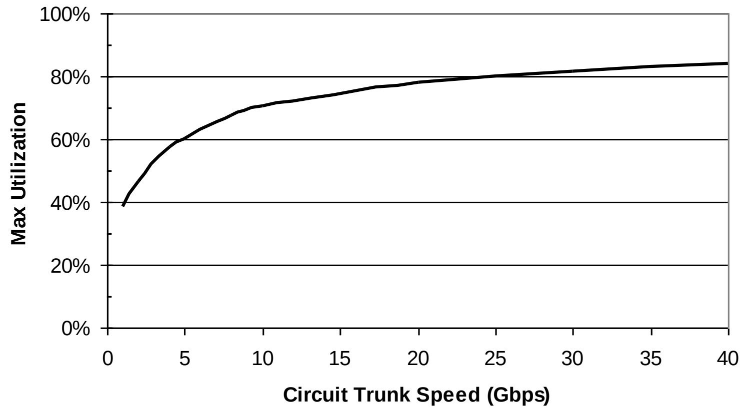

Figure 2. Maximum Utilization versus Circuit Trunk Speed for High Speed Internet Access

3. Benefits of Cooperation between Packet and Transport Networks

If the transport network can communicate performance characteristics (e.g., latency due to propagation delay) to a packet network dynamically, then the packet network can make better informed decisions about how to forward packets such that quality objectives are met while also optimizing efficiency.

If the cost of a transport node supporting packet switching, for example, using the MPLS Transport Profile (MPLS-TP) is much less than that of an MPLS only Label Switch Router (LSR), then the overall network cost can be substantially reduced as shown in Figure 2.

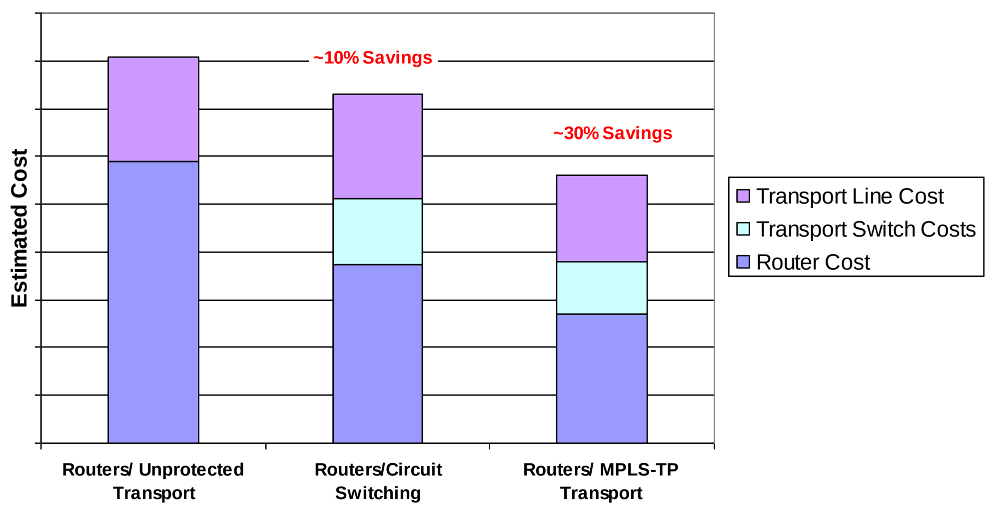

Figure 2. Relative Cost of Networks with and without Cooperation between Packet and Transport

The leftmost bar shows the relative cost for routers connected over unprotected transport where the routers perform restoration in response to failures. The middle bar shows the relative cost when packet traffic is carried over routers connected via a circuit switched transport network which performs restoration. Approximately 10% savings occurs in this case because the cost of circuit switching is substantially less than that of router ports. These models did not include the impact on maximum utilization due to circuit trunk speed described previously since all circuits were chosen to be of the same speed. The rightmost bar shows the relative cost when the transport network performs packet switching (e.g., MPLS-TP). The increased savings of approximately 30% results because the number of more expensive MPLS LSR nodes which an MPLS Label Switched Path (LSP) traverses is reduced, for example, in the case of a full mesh the LSP traverses only two MPLS LSRs.

In parts of the network where the aggregate circuit switched and packet switched demand is less than that of a single wavelength, the ability of the transport network to allocate a fraction of capacity in ODU-0 increments (approximately 1.25 Gbps) allows more efficient wavelength utilization and sharing of capacity between TDM and packet demands as compared with allocating separate wavelengths to the circuit and packet networks.

If on-demand circuit switched services are provided, then efficiency is increased (or performance in terms of blocking probability is improved) as the pool of available capacity increases. A pool of capacity cooperatively shared between on-demand circuit switched services and packet switch interconnection, subject to policy controls that limit the total capacity usable by either network, could results in improved efficiency for both services.

4. Potential Methods of Cooperation between Packet and Transport Networks

When a packet network uses dynamically signaled MPLS in cooperation with a transport network, for example, using the dynamic subset of MPLS-TP [3] standards being jointly developed by the IETF and the ITU-T that aim to make the transport network packet switching capable by defining a subset of MPLS while adding some traditional transport OAM capabilities. In this mode of operation, LSRs can create LSPs that directly connect sites with a large community of interest. This LSP connectivity may be a full mesh if traffic demands warrant such connectivity.

As described previously, when packet and transport networks perform restoration independently, the possibility that the performance (in particular, latency) observed by the packet network can be adversely impacted. There are at least two methods for solving this problem. The first is for the packet network to measure latency provided by underlying transport network(s), and then react based upon these measurements. This method requires no cooperation between packet and transport networks, but is subject to potential misinterpretation of other factors (e.g., congestion, packet routing) as a change in long term latency, which can lead to oscillations in routing. The second requires cooperation between a packet and transport network. In one method, the packet network could request of the transport network via automatic signaling that a (virtual) circuit have latency no greater than a certain value. In another method, the transport network could communicate to a packet network in response to an automatically signaled (virtual) circuit the estimated latency of the established (virtual) circuit. If the transport network performed a restoration action which altered latency, this could also be communicated via this automatic signaling and the packet network could respond accordingly.

Generalized MPLS (GMPLS) signaling [4] can be used to signal Lambda or Circuit services by end users, other networks, or a proxy, and also support dynamic signaling by routers to establish circuits of the required speed and performance characteristics (e.g., latency). As mentioned previously, since packet switching over a set of circuit switched trunks has lower maximum utilization, the speed of the circuits may need to be adjusted relatively frequently by the routers in response to changing demands in order to achieve acceptable efficiency.

When a transport network uses circuits of a rate less than the interface speed used to interconnect routers, there must be a means to match the average rate of the router packet transmission to that of the circuit. Hierarchical scheduling [5] is a commonly used technique in packet access networks (e.g., Digital Subscriber Line (DSL), FTTx, and Enterprise network access) is to perform this sort of rate matching and still preserve multiple QoS classes. If the router signals the circuit speed dynamically via GMPLS, then the router can also adjust the rate of the hierarchical scheduler and the rates of the QoS class queue schedulers accordingly. Without such a cooperative mechanism, performance degradations such as excessive loss and/or delay variation could occur.

5. Conclusions

This paper described how performance degradations and inefficiencies can occur when packet and transport networks operate autonomously. An example of the economy of scale in statistical multiplexing efficiency was given, demonstrating the region where circuit switching can be used cost effectively. One important solution is for the packet and transport networks to cooperate via signaling and routing so that less expensive circuit or packet switching functions implemented in the transport equipment reduce overall network cost. Furthermore, an important performance parameter, latency can be better controlled if packet and transport networks cooperate.

6. References

[1] R. Guerin, "Equivalent Capacity and Bandwidth Allocation," IEEE JSAC, Sept 1991.
[2] C. Villamizar, D.McDysan, Editors, "Requirements for MPLS Over a Composite Link," IETF, Work in Progrss.
[3] B. Niven-Jenkins, D. Brungard,, M. Betts, Editors, "Requirements of an MPLS Transport Profile," IETF, RFC 5654, September 2009.
[4] D. Papadimitriou, J. Drake, J. Ash, A. Farrel, L. Ong, "Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON)," IETF, RFC 4139, July 2005.
[5] DSL Forum, "DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services," TR-059, September 2003, http://www.broadband-forum.org/technical/download/TR-059.pdf.