TransPacket white paper

Saving costs by increasing capacity utilization of optical networks

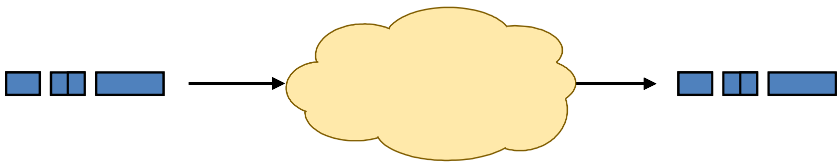

04.01.2013

Intelligent traffic injection in networks

Executive summary

Capacity utilization of optical networks is typically very low. IP/MPLS and Ethernet packet optical networks are replacing SDH based systems for long distance and metro networks to improve flexibility and bandwidth utilization. However, it is desirable to keep the unique QoS and timing characteristics of the circuit systems. Keeping the load of packet systems at a moderate level enables high QoS, but still without reaching the true circuit QoS characteristics. Fusion Networking is a new technology enabling both packet network capacity utilization and true circuit properties. This enables operators and enterprises to improve network resource utilization, saving cost and increasing flexibility and may result in doubling the revenue.

Saving costs in fiber systems by increasing capacity utilization

Load and utilization of wavelengths in WDM networks

Typically, wavelengths in WDM packet networks are far from being fully utilized. The traffic load follows statistical variations during the time of day, and the peak load in a typical WDM network never reaches a point where the wavelengths capacity is saturated. As an example, figure 1 shows a example measurement of the capacity utilization during 24 hours on a 10 Gigabit Ethernet link between Trondheim and Oslo in the network of the carrier UNINETT [1]. The figure shows that, the capacity utilization peaks at approximately a load of 0.35 (3.5 Gb/s).

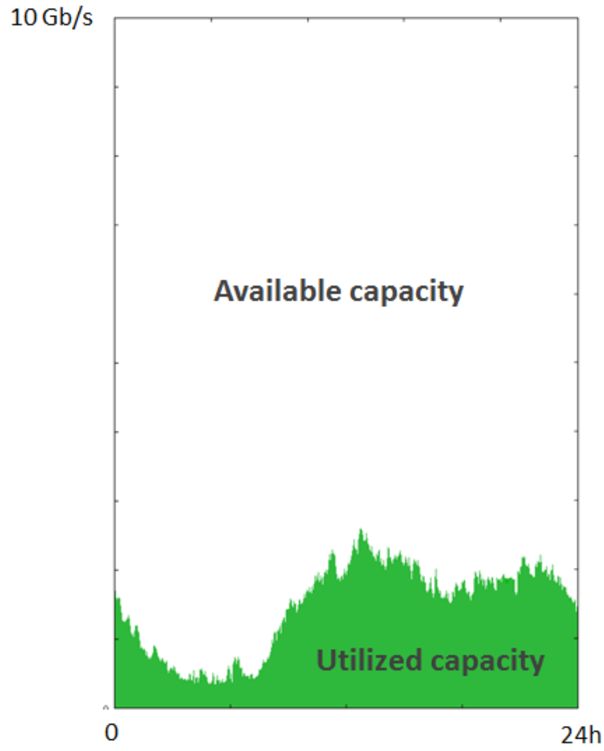

Figure 1: Capacity utilization of a 10GE wavelength during 24 hours on the Trondheim-Oslo link in the network of UNINETT in Norway.

There are several reasons why utilization of wavelengths is held low in carriers networks:

1) Carriers may have a policy to invest at 50 % fill rate to avoid potential overload and corresponding impact on Quality of Service (QoS).
2) A large amount of the wavelength capacity may be reserved for protection switching purposes, handling added traffic in case of a failure in a redundant link.
3) Ethernet capacity comes in chunks of gigabit or 10 gigabit. When a customer´s bandwidth demand rises beyond one gigabit, renting several gigabit Ethernet connections may have a higher cost than renting a 10 gigabit connection. Jumping to a 10 gigabit line or wavelength will then result in low wavelength capacity utilization.

As figure 2 illustrates, the average capacity utilization of an optical link will typically be far less than 25 %, given a peak utilization at the 35 % as experienced during the logged traffic pattern during 24 hours in figure 1.

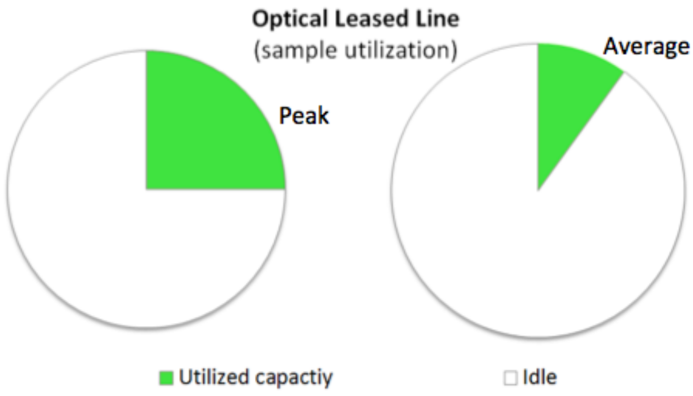

Figure 2: Illustration of typical capacity utilization in a long-distance optical wavelength. Most of the capacity is always vacant, ensuring a high QoS level but heavily reducing potential revenue.

If vacant capacity can be utilized without any impact on the services, it is possible for carriers to more than double the utilization of their wavelengths, postponing investments in leased lines or WDM infrastructure. This white-paper explains how to increase capacity utilization of wavelengths without any impact to current services, and how carriers can benefit in terms of cost-savings in their network.

Exploiting idle wavelength-capacity

When a customer rents a wavelength from a carrier, the Service Level Agreement (SLA) typically dictates that the wavelength is fully reserved for the purpose of the customer. Inserting an Ethernet switch or IP-router will impact the performance of the customer traffic in the wavelength. Offering a wavelength quality link using a TDM system like SONET/SDH or OTN is an option, but unlike Ethernet switches and IP-routers, statistical multiplexing allowing efficient bandwidth utilization is not available. Hence, the carrier is inhibited from utilizing any of the idle capacity of the wavelength.

The fusion networking principle solves this problem. By using an Intelligent Traffic Injection (ITI) mechanism, idle capacity can be utilized without any impact to the traffic in a 10 Gigabit/s wavelength or in 1 Gigabit/s sub-wavelengths. By carefully monitoring the packet-traffic in the wavelength/sub-wavelengths, gaps between packets representing idle capacity is found. The ITI mechanism allows traffic to be injected in these vacant gaps without any impact on the packets that are already following the wavelength or sub-wavelength. Hence, at the customer side the traffic injection is not measurable, and performance of the link equals a transponder based link: No packet loss, low latency and ultra-low packet-jitter.

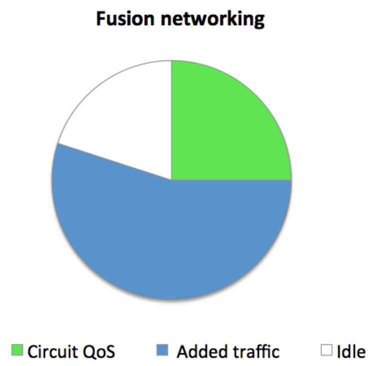

Figure 3, increased capacity utilization by adding packet traffic onto a link without impact on existing traffic.

Figure 3 illustrates how additional packet traffic can be injected on a link or wavelength. The link utilization may be increased to 80 %, injecting 55 % of additional traffic. The traffic being injected in the wavelength has a lower priority than the customer traffic following the wavelength. The traffic following the wavelength is called Guaranteed Service Transport (GST) traffic and receives absolute guarantees. The injected traffic is Statistically Multiplexed (SM) onto the wavelength traffic, and receives statistical guarantees. Hence, the guarantees for the SM traffic yields as for a conventional packet switched network: Throughput depends on the traffic in the network and the oversubscription rate if oversubscription is applied. By monitoring the 10 gigabit/s Ethernet traffic load, the amount of idle capacity can be found. Alternatively, if the 10 Gigabit/s Ethernet wavelength is subscribed with a maximum of e.g. five Gigabit Ethernet sub-wavelengths, a leftover capacity ranging from a minimum of five up to 10 Gigabit/s will be available for traffic injection. When injecting additional traffic on the wavelength, dimensioning considerations takes into account how much traffic that is available and which oversubscription rate being acceptable.

Double the revenue

If a wavelength runs at e.g. only 35 % peak-utilization, four to five additional gigabit Ethernet lines can be added, still guaranteeing performance of zero packet loss and latency in the lower millisecond range. Long distance transport between cities or even across continents is generally far more expensive than local transport within cities, making efficient bandwidth utilization on long-distance links especially attractive. In this white-paper we illustrate how the flexibility and increased capacity utilization may enable revenue stemming from individual links to be doubled.

HOW TO SAVE COSTS IN FIBRE NETWORKS

Higher cost-efficiency on 10 Gigabit Ethernet wavelengths using H1 can be achieved through two main business models, called Sub-wavelength and Traffic injection business model, explained below.

1) **Sub-wavelength business model:** Splitting a 10 Gigabit/s wavelength into Gigabit/s sub-wavelengths, offering the sub-wavelengths to customers and dynamically using any leftover capacity in the 10 Gigabit/s wavelength for internal transport needs, as illustrated in figure 4.

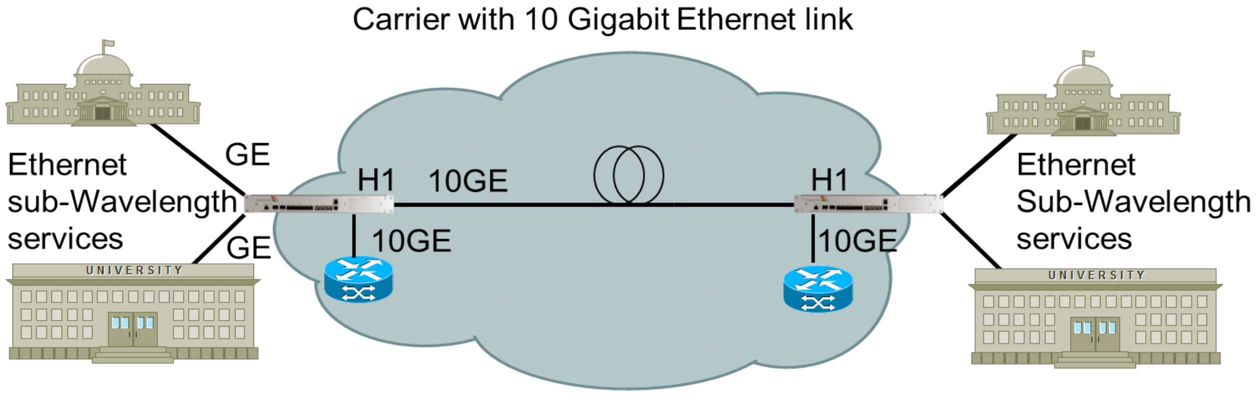

Figure 4, Using TransPacket H1, a 10 Gigabit Ethernet (10GE) wavelength is split into single Gigabit Ethernet sub-wavelengths and offered to customers as Gigabit Ethernet wavelengths with wavelength QoS guarantees. Any leftover capacity on the 10 GE is available for the carriers own transport needs.

In this business model, the carrier can sell several sub-wavelength 1 GE connections with wavelength QoS guarantees to its customers. At the same time, the carrier can leverage unused capacity between routers connecting to the H1 using 10GE. The sum of revenue from several guaranteed 1GE connections typically outnumber the revenue possible from one 10GE, and in addition any unused traffic capacity is available for communication between the routers or switches of the carrier. As an dimensioning example, If the carrier rents out five 1 GE connections to its customers, each of the customers loading their links to a mean of 50 %, the available capacity on the 10GE link will be: Minimum 5 Gb/s, maximum 10 Gb/s and a mean of 7.5 Gb/s. Comparing with a SDH/SONET or OTN system, the maximum capacity available would be 5 Gb/s.

2) **Traffic injection business model:** Offering a 10 Gigabit/s wavelength to a customer, monitoring the traffic load of the wavelength and use any leftover capacity for internal transport needs, as illustrated in figure 5.

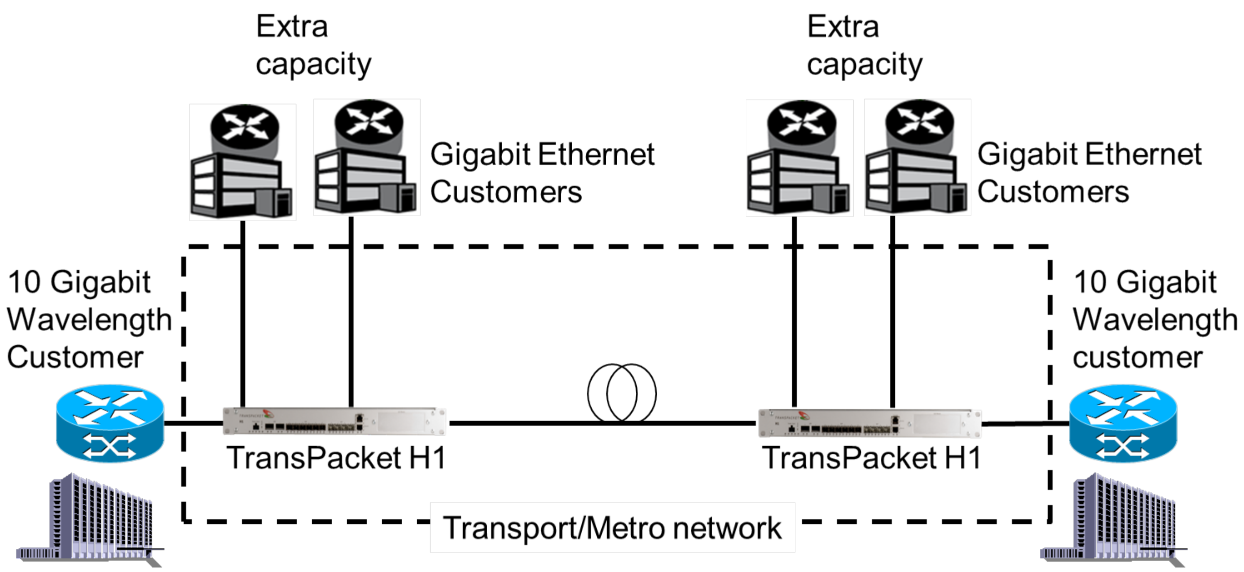

Figure 5, a 10 Gigabit Ethernet wavelength offered to a customer typically has moderate bandwidth utilization. By monitoring the load on the wavelength, any vacant capacity may be offered as gigabit Ethernet services to customers requiring a packet switched Ethernet link service level.

When renting out a 10GE wavelength, or link with wavelength quality, exploiting vacant capacity has so far not been an available option. By using the FUSION networking principle, any leftover capacity becomes available for traffic injection. Using H1, the leftover capacity is available at the ten GE interfaces. As an dimensioning example, if the 10GE customer loads up his link to a peak-load of 50% and a mean load of 25 %, there is still a minimum 5 Gb/s and a mean of 7.5 Gb/s available on the link. This can be utilized either by connecting routers/switches of the carrier, or sold as e.g. GE links to the carriers customers. The SLA of the GE links should then be similar to the oversubscribed GE service typically offered to business customers not requiring a private line.

THE ADDED VALUE OF FUSION NETWORKING

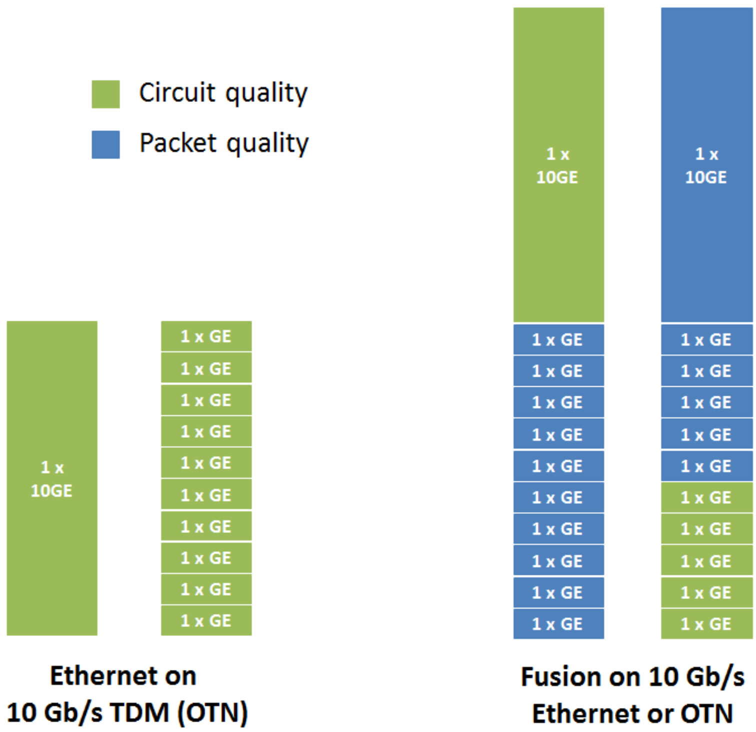

Figure 6, Fusion networking offers the same quality as for circuit switched TDM, but with the added value of offering more interfaces, enabling the same capacity utilization as for packet switched networks.

Because network capacity utilization typically is tripled, flexibility increased, and circuit QoS is made available, revenue stemming from individual links may be doubled. Figure 6 illustrates the increased capacity utilization and flexibility through Fusion networking. Using TransPacket FUSION H1, e.g. a 10 Gigabit/s SDH or OTN circuit can be loaded through either a 10 Gigabit/s Ethernet interface or ten Ethernet interfaces of 1 Gigabit/s. The FUSION network offers a different level of flexibility but with the same guarantees as circuit switching. A 10 Gigabit/s Ethernet or OTN circuit can be loaded using a variety of interface combinations. As an example, a 10 Gigabit/s interface can offer circuit guarantees while additional ten interfaces of 1 Gigabit/s offers additional traffic injection. This is compatible with businesscase 1) given above. Another example illustrating business-case 2) is as follows: Rather than offering ten 1 Gigabit/s interfaces as in circuit switching, FUSION can e.g. offer five guaranteed 1 Gigabit/s Ethernet interfaces and offer traffic injection through an additional 10 Gigabit/s Ethernet interface and five Gigabit/s Ethernet interfaces. The traffic being injected will then have bandwidth guarantees with a minimum of 5 Gigabit/s and a maximum of 10 Gigabit/s available.