

Internet Traffic Characteristics

- *Bursty Internet Traffic*
 - Statistical aggregation of the bursty data leads to the efficiency of the Internet.
- *Large Variation in Source Bandwidth*
 - 10BaseT (10Mb/s), 100BaseT(100Mb/s), GbE(Gb/s)
 - Congestion control is needed, not just to deal with gross overload, but also slows down sources to an economically sustainable operating point.
- *Unpredictable Traffic*

Ref: David D. Clark, "Fiber-based metropolitan access networks for Internet traffic", OFC'2000, paper WD1.

1

How to take care of the Bursty IP traffic in Optical Networks

- **Q:** How optical networks, which have no buffering, could carry the bursty sorts of traffic that the Internet generates.
- **Ans:** Since no optical buffer is available, electronic buffer is needed to aggregate traffic to provide somewhat smooth total traffic sources before loaded onto fibers.

2

IP-over-WDM

- A more simple and low-cost overall network design is achievable if *the lower layers of the system have been designed specifically to carry IP traffic.*
- *need to simplify or eliminate the layers that may have been running under IP, e.g. Frame Relay, ATM, SONET, and replace them selectively with a set of services that explicitly meet the needs of the Internet protocols.*
- The set of problems that these services must solve include *bandwidth management, restoration, and support for packet routing.*

3

Can WAN model be used in access networks?

- The approach required in an access network to adopt IP-over-WDM may be quite different from that in a wide area network.
 - For example, the topology of a distribution network is usually a tree, in physical form.
 - The methods used for restoration of a SONET ring do not directly map onto a tree topology.
- The tradeoff between approaches at different layers becomes part of the overall design process.
- For access network, which is cheaper—*scheduling mechanisms* (like WAN) or *excess capacity* (like Ethernet)?

4

Restoration

- **Q** : How to implement Restoration on IP-over-WDM?
- In case of a link failure, one can reallocate the remaining capacity among all the paths, the flows will adjust their sending rates to compensate for the loss of bandwidth.
- This approach, however, needs to reassign active fiber capacity in a short time frame (*reconfiguration time*), which is a different problem than the activation of idle restoration paths.

5

New Protocols?

- Protocols such as TCP contain flow and congestion control algorithms that slow the transfer down when necessary.
- The time constants of the current implementations may not perform well at multi-gigabit speeds.
- If IP-level routing shifts the packets onto a new dedicated path (WDM channel),
 - the TCP should adapt quickly to the increased capacity of the link, and
 - the application transfer should speed up to the capacity of the optical path without resetting of any of the software layers.
- If the dedicated optical path is then discontinued,
 - the Internet protocols will fall back to the normal slower path for IP data.

6

Summary

- In general, operation of Internet protocols at high speeds will require some redesign of rate control and congestion algorithms and some tweaking of implementations.
- TCP is a transport protocol that has to adapt to both fast and slow links.
- With the advent of real gigabit long distance networks, these algorithms will have to adapt over a dynamic range of at least 10^5 in speed and in distance.
- The development of very dynamic control algorithms is part of current Internet research.

7

Traffic model for USA long-distance optical network

Traffic model for USA long-distance optical network

Anurag Dvivedi and Richard E. Wagner,

Optical Network Research, SP-AR-02-277U, Corning Incorporated, Corning, NY-14851
Phone: 607-974-4237, Fax: 607-974-9271, Dvivedia@Corning.com

Abstract

Voice, transaction data, and Internet traffic models are developed and used to estimate the traffic demand and growth between major US cities. Shortest path routing is used to estimate the capacity required on individual links, as well as traffic granularity and the fraction of traffic dropped at each city.

OFC'2000 Paper TuK1

8

Traffic models for USA long distance Optical Networks

Traffic models can be used to estimate network parameters that are very useful to determine equipment sizing, functionality, and requirements.

The parameters of interest include

- *the capacity needed in individual links in the network,*
- *the granularity of the traffic on these links,*
- *the fraction of traffic dropped at any node, and*
- *the distribution of link lengths in the network.*

9

Traffic models for USA long distance Optical Networks (cont.)

- A comprehensive traffic model of the USA long-distance network is developed by segregating the traffic into three application areas (*voice, transaction data, and internet data*) and determining the traffic demands and growth rate of each application.
- Subsequent traffic routing and network analysis yields the statistics of the parameters mentioned above directly from the models.

10

Traffic Model

$$\left. \begin{aligned} \text{Voice traffic}(i,j) &= K_V \cdot P_i \cdot P_j / D_{ij} \\ \text{Transaction data traffic}(i,j) &= K_T \cdot E_i \cdot E_j / \sqrt{D_{ij}} \\ \text{Internet traffic}(i,j) &= K_I \cdot H_i \cdot H_j \end{aligned} \right\} \text{Eq.1}$$

where the traffic between major cities i and j depends on the total population P_i , the non-production business employees E_i , and the number of internet hosts H_i in each city, as well as the distance D_{ij} between the two cities of interest.

The proportionality constants that define the traffic levels assume

- 14 minutes of long-distance voice traffic per person per day,
- 5 minutes of transaction modem use per non-production employee per day, and
- 25 minutes of continuous modem access to the internet per host day.

11

US Long-Haul Networks

- Total traffic demand of 214 Gb/s in 1998 for the USA, with a breakdown of 77%, 11% and 12% for voice, transaction data, and internet applications, respectively. (Expected in 2000 or 2001, voice will be less than 50%)



Figure 1: US long-haul network showing the individual link capacity and average granularity

12

Traffic growth rates

Table 1: The methodology for estimation of traffic growth rates

	Voice Traffic $K_v \cdot P_v \cdot P_f / D_v$	Transaction data Traffic $K_t \cdot E_t \cdot E_f / V D_t$	Internet Traffic $K_i \cdot H_i \cdot H_f$
Growth Factors	P @ 1.7%/yr	E @ 4%/yr	H @ 44%/yr
	K_v @ 5%/yr	K_t @ 24%/yr	K_i @ 24%/yr
Total Growth	8%/yr	34%/yr	157%/yr

- The composite aggregate USA long-distance traffic is projected to grow from 275 Gb/s in 1999 to 8 Tb/s by the year 2004, five years from now.

13

Traffic growth rates

- Taking into account a peak load to average traffic ratio of 5, allowing another 60% excess capacity for restoration and protection, and accounting for the finite equipment transmission rates of 2.5 Gb/s and 10 Gb/s, the installed capacity needed to support 8 Tb/s of average traffic is about 73 Tb/s.

14

Link Length and Connection Distance

- The statistics of link lengths and connection lengths in the network are shown in Figures 2 and 3 respectively.




Figure 2: Distribution of link distances





Figure 3: Distribution of connection distances

15

Distribution of Add/drop traffic


- By routing the individual traffic demands generated by the Eq (1), and collecting all of the individual traffic demands that each link in the network supports, it is possible to determine other network parameters, such as *total link traffic* and *link granularity* (both illustrated in Figure 1) and *the fraction of traffic dropped at each city*.



16

Projection of 10 Gb/s optical channel need

- By applying the growth rates to the three traffic applications and aggregating the resulting traffic demands, it is possible to project the future needs for 10 Gb/s optical channels on each link. (for link capacities needed in the year 2004).



17

Conclusion

- Developed a comprehensive traffic model for USA long-distance traffic that allows traffic growth projections and subsequent network requirements to be analyzed.
- The international components of these applications are not yet included in the model and these will increase the demands by 10% or so.

References:

1. Sinclair, M. C., Electronics Letters, Sep. 1, 1994, Vol 30, No. 18, pp. 1468-1470
2. Population data from US Census Bureau, URL: www.census.gov
3. Growth chart for non-production and production employees published by the census bureau, http://www.census.gov/epcd/cbp/images/cbp_hist.gif
4. Internet hosts statistics, <http://www.mids.org/growth/internet/index.html>
5. The Size and Growth rate of the Internet, Kerry Coffman and Andrew Odlyzko, http://www.firstmonday.dk/issues/issue3_10/coffman/index.html
6. The Statistics of Communications Common Carriers, FCC, <http://www.fcc.gov/ccb/stats>
7. GVU WWW user survey, http://www.cc.gatech.edu/gvu/user_surveys/

18

Reconfiguration in IP Over WDM Access Networks

Reconfiguration in IP Over WDM Access Networks

Jennifer Yates, Gisli Hjálmtýsson and Albert Greenberg

AT&T Labs – Research, AT&T Shannon Laboratories, Florham Park, NJ 07932-0971, USA
Phone: (973) 340 7036; Fax: (973) 340 8050; {jyates, gisli, albert}@research.att.com

This paper examines a high-speed IP over WDM access network. The performance benefits of dynamically allocating wavelengths to access routers are examined and compared with static wavelength allocation. It is shown that dynamic reconfiguration can significantly improve network performance.

OFC'2000 Paper TuK4

19

NGI ONRAMP

- The *NGI ONRAMP* consortium is focusing on providing high speed optical access to businesses.
- *ONRAMP* networks are metropolitan or small regional networks, consisting of a single feeder ring network with approximately 10 access nodes.
- These networks are to be realized using *IP over WDM*, with no intelligent networking layer (e.g. ATM) in between.
- **The aim :**
to build thousands of regional ONRAMPs, each being low cost and easy to provision and manage.

20

NGI ONRAMP

- **The Assumption:** *negligible local traffic between access nodes*

Optical *lightpaths* (wavelength) are established to transport IP traffic from an access node to the gateway without passing through intermediate IP routers.

These optical lightpaths may be *statically* allocated or may be *dynamically* reconfigured to either follow long-term shifts in offered loads, or more rapid transient load fluctuations.

21

ONRAMP Ring Feeder Network

- These access nodes are used to connect high-speed customers and distribution networks to the feeder ring.
- The feeder ring is also connected to a backbone network via a gateway node, or a group of co-located gateways.
- Each access node consists of an *IP router* and an *optical cross-connect (OXC)* (or an OADM).




Fig. 1. Example uni-directional ring feeder network with N access nodes and a single gateway. Each access node consists of an optical cross-connect (OXC) and an IP router. ONRAMP consists of all elements below the gateway.

22

Wavelength Assignment

- Assume # of wavelength > # of access node
→ each node can get at least one wavelength
- The gateway node allocates the wavelengths to the different access nodes.
- simplest assignment : the wavelengths are allocated statically according to traffic forecasts and provisioning rules.
- More efficient way is to assign the wavelength dynamically.
→ need tunable transmitter and receiver in each IP router

23

IP Flow

- An *IP flow* consists of a set of IP packets with IP header fields.
- IP flows arriving at an access node are assigned to the wavelength at that node with the *smallest number of flows*.
- Flows are not divided across multiple wavelengths.
- Multiple flows are then aggregated onto each wavelength using a fair scheduler.

For example, if the capacity of a single wavelength is C , and there are M flows currently being transported on this wavelength, then each flow receives a bandwidth of C/M .
- Heavy tailed *Pareto distribution* is used to describe the amount of data transmitted within a flow and with flows arriving according to a *Poisson distribution*.

24

“Slow Down”

- The performance of the WDM feeder network was measured using the *slow down* experienced by flows in the system.
- The *slow down* is defined to be the ratio of *the time taken to transmit a flow* to *the time required to transmit the flow if it had dedicated access to an entire wavelength*.
- Ideally, flows experience a small slow down (close to 1) with very small probabilities of flows experiencing large slow downs.

25

wavelength reconfiguration algorithm

- Keeping all wavelengths in the system evenly loaded will minimize the mean slow down.
→ The aim of the wavelength reconfiguration algorithm is to *keep the load on each wavelength approximately equal*.
- The mean load on the wavelengths at a node j is given by n_j/W_j , where n_j is the number of flows at the node, and W_j is the number of currently assigned wavelengths.

26

wavelength reconfiguration algorithm

- Wavelengths are moved from node k with the minimum n_k/W_k to node j with the maximum n_j/W_j , subject to the constraint that :
 - all nodes should have at least one wavelength allocated, and
 - $n_j/(W_j+1) + n_k/(W_k-1) < n_j/W_j + n_k/W_k$
(i.e. the network is more evenly balanced as a result of the wavelength move).
- To model signaling and laser tuning times we introduce a delay, D , after the decision to move a wavelength has been made until it takes effect.

27

The probability of various slow down

- We thus examine the slow down distribution to evaluate network performance.




Fig. 2. Slow down distribution for a 10 node network with 16 wavelengths.

- Traffic is uniformly distributed amongst all nodes with a per wavelength load of 0.7 and with Pareto flow holding times, with $\alpha = 1.2$, truncated at 10^5 .
- In static mode, each node is assigned with two wavelength.

28

Results

- The fig. shows that the slow down is substantially reduced by wavelength reconfiguration.
- Mean slow down is reduced from 1.95 to 1.008.
- More significantly the tail of the slow down distribution is dramatically reduced.
- In comparison, with full reconfiguration each router can have up to 9 wavelengths (maximum of 8 wavelengths can be moved as each router always has at least one wavelength allocated), and thus needs 9 ports.
- With the 4 port routers we see only small performance degradation as compared with the 9 port routers. This reduction in router size is significant.

29

Standards for IP-over-WDM

- Convergence of data-communications and telecommunications networks
→ convergence layer: IP (?)
- **Q:** How to develop IP-over WDM to facilitate cost-effective delivery of the full range of voice and data services over a unified network infrastructure?
- **Ans:** Two primary approaches toward meeting this objective are
 - *optical internetworking* and
 - *optical-data networking*.

Ref: Lightwave, March 2000

30

Optical-Internetworking agreements

- *Optical internetworking* is under consideration within the optical-internetworking forum (OIF).
- It represents a *data-optimized* network infrastructure in which switches and routers have integrated with optical interfaces and are directly connected by fiber or optical-network elements, such as dense wavelength-division multiplexers (DWDM).
- Here, *IP over WDM* means IP packets mapped into Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH), coupled with SONET/SDH-based point-to-point DWDM systems.
- The OIF is working toward optical-internetworking agreements, which it expects to publish in the 2000/2001 timeframe.

31

Optical-data networking

- Optical-data networking is under consideration in Committee T1X1 and the ITU-T.
- It represents carriage of data over a service-transparent infrastructure that provides optical transport networking (OTN) features, including optical-channel (wavelength) routing, fault and performance monitoring, and protection and restoration.
- Here, *IP over WDM* means packets directly mapped-via encapsulation-into an optical channel, where any client signal is accommodated, regardless of its protocol format and bit rate.

32

Optical-data networking (cont.)

- This procedure involves usage of a *digital wrapper* "around" the optical-channel client to support channel-associated operations, administration, maintenance, and provisioning functionality.
- The optical channel may then be networked and managed without being constrained to follow point-to-point links coupled with router topology.