FLAMINGO: A Packet-switched IP-over-WDM All-optical MAN

D. Dey (1), A. van Bochove (2), A. Koonen (3), D. Geuzebroek (4), M. Salvador (5)
(1) Faculty of Electrical Engineering, University of Twente, Netherlands, dey@cs.utwente.nl
(2) KPN Royal Dutch Telecom, Netherlands, a.c.vanbochove@kpn.com
(3) COBRA Institute, Eindhoven University of Technology, Netherlands, a.m.j.koonen@tue.nl
(4) Faculty of Electrical Engineering, University of Twente, Netherlands, d.h.geuzebroek@student.utwente.nl
(5) Faculty of Computer Science, University of Twente, Netherlands, salvador@cs.utwente.nl

Abstract: We describe network architecture of an all-optical wavelength-and-time-slotted MAN. Key aspects of this architecture include all-optical packet switching, the ability to put IP packets directly over WDM and the possibility of interfacing with any heterogeneous network.

1. Introduction

Metropolitan Area Networking is a lucrative market opportunity. The advantages of successful deployment of WDM accompanied by packet switching in the MAN are incontrovertible. Optical time-slotting may be considered as one of the techniques useful in realising a packet-switched network (see e.g. /1/ and /2/).

In this paper we present the architecture of a WDM Metropolitan Area Network (MAN) that goes beyond these studies by incorporating aspects such as all-optical packet switching, IP over WDM, and interfacing with any heterogeneous network, a feature very desirable from the operator's point of view.

2. The MAN

The FLAMINGO[1] network architecture is as shown in figure 1. It is a slotted WDM network laid out as a fiber ring. Access to the MAN is via the access points (APs) as shown in figure 1. The MAN is thus connected to a Wide Area Network (WAN) and heterogeneous types of LANs via the APs.

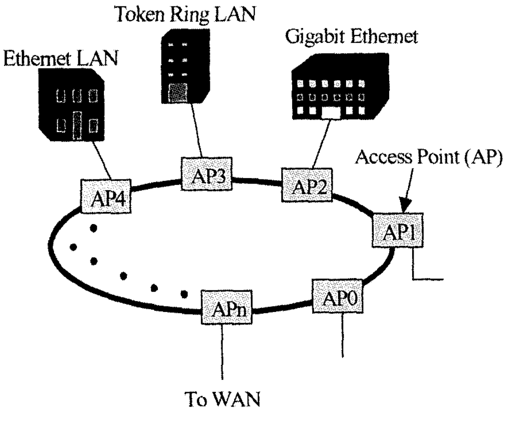

Figure 1: The FLAMINGO MAN

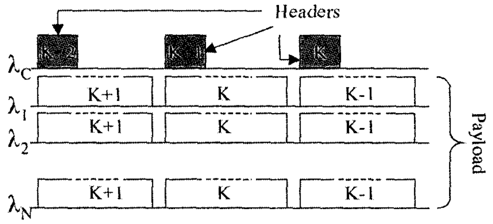

Figure 2: Synchronised Payload and Header slots

Bandwidth of each WDM channel is divided in time-domain into equal and constant number of slots that circulate around the ring. One of the WDM channels is reserved for carrying control information and the header of all the payload channels. The remaining channels carry the payload. This is shown in figure 2. Each AP on the ring is able to transmit and receive at any wavelength. A limited case in which an AP is able to transmit on one wavelength and receive at any wavelength has been analysed in /3/ and /4/. The header slot and the payload slots are synchronised such that the header slot arrives a notch ahead of the payload slots as shown in figure 2. The header slot informs the AP if any payload slot on any wavelength is empty for transmitting data or if any payload-slot on any wavelength needs to be dropped.

To enable fast and efficient header processing we consider the usage of labels in the header, each label corresponding to a payload slot on a WDM channel.

3. Access Point (AP)

The access point is as shown in figure 3. The incoming data is first put through a tuneable λ-drop to separate the header channel from the payload channels. The tunability protects the network against a failure in the existing header channel, an event that will bring down the whole network if unprotected. The payload slots are then delayed in a fiber loop. This delay, induced by the Header Processor unit (HPU in figure 3), is limited by the header slot size at a given line speed. The payload channels are then de-multiplexed. If a payload slot is not empty, then, depending upon its destination, it may either be allowed to bypass an AP or be dropped at an AP. The AP is able to add data to a payload slot if it is empty. One of the major drawbacks of such a network is dispersion, as a result of which slots on various channels may lose synchronisation. This is however taken care of in each AP. Every time an AP

---

[1] FLAMINGO (Flexible Multiwavelength Optical Local Access Network Supporting Multimedia Broadband Services) is a research project supported by the Dutch Technology Foundation, STW.

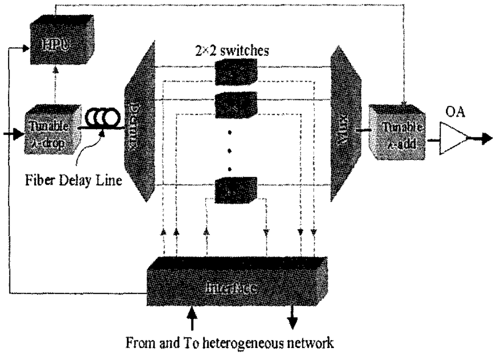

Figure 3: Access Point (AP)

generates data it sends a regenerated signal whose time of transmission can be controlled within a desired level of accuracy.

Interfaces are required at the APs to support incoming SDH frames from the WAN and Ethernet and other types of frames from the LANs and vice versa. This is done at the AP interfaces to these heterogeneous networks (shown in figure 3). The interfaces tear down these frames into IP packets, read their IP headers at the HPU to locate their destination APs and then put them onto incoming empty slots.

To synchronise all the HPUs in the ring we propose using the Embedded Clock Transport technique, in which the transmitter frequency multiplexes its clock with the baseband data and modulates the optical carrier with the composite signal. Since the network is a ring complete bit-synchronisation is impossible. Sufficient gap must therefore be maintained between the header and the payload to deal with the asynchrony.

5. Performance Results

The above network was simulated in C++. The network consisted of a 35-kilometer unidirectional fiber ring with 10 APs (nodes) evenly spaced. The system was assumed to transmit 4 wavelengths at 2.5 Gbps per wavelength for a total fiber capacity of 10 Gbps. One of these wavelengths was reserved for carrying control information and the headers while the others were reserved for carrying the payload.

The network was simulated with 3 different slot-sizes, namely, 44, 552, and 1500 bytes circulating around the ring simultaneously one after the other. The percentage of slots of each type in the ring equalled 50%, 30% and 20%. Such IP-packet size distributions have been reported in /5/ and /6/. The packet arrival distribution was assumed to be a homogeneous Poisson process with rate μ (a constant for all the APs in the ring) and the graphs obtained are shown in figures 4 and 5.

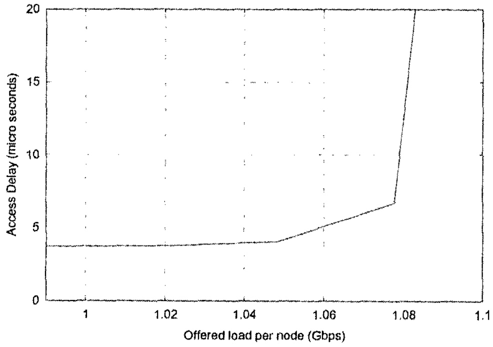

Figure 4: Offered load per node versus access delay

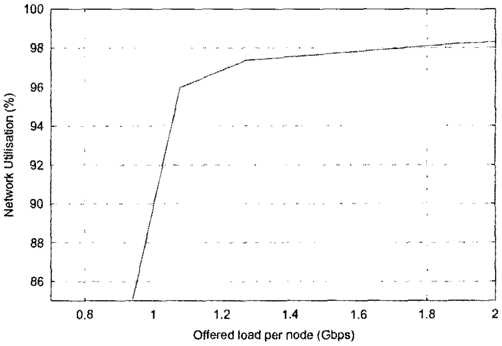

Figure 5: Network utilisation versus offered load per node

In figure 4, access delay implies the average amount of time an IP packet has to wait from the instant it enters a buffer at an AP till an empty slot picks it up. In figure 5, network utilisation implies the percentage of network resources (in this case, bandwidth of the total WDM system) used up to support the offered load.

6. Conclusions

Together with figure 5, we conclude from figure 4 that relatively high throughputs at an AP can be achieved in the FLAMINGO architecture as a result of the high degree of network utilisation. The ability to put IP-packets directly over WDM channels enables bit-rate and protocol transparencies within the MAN and helps an operator interface with almost any heterogeneous type of network.

References

/1/ I. Chlamtac, et al., "Scalable WDM Access Network Architecture based on photonic slot routing", IEEE/ACM trans. on Networking, vol.7, no. 1, (1999), pp. 1-9.

/2/ C. -S. Kang, et al., "A broadband ring network: Multichannel optical slotted ring," Computer Networks and ISDN systems, vol. 27, (1995), pp. 1377-1398.

/3/ D. Dey, et al. "Network Architecture of a Packet-switched WDM LAN/MAN," Proc. of IEEE/LEOS Benelux Chapter, (2000), pp. 256-259.

/4/ M.R. Salvador, et al. "An all-optical WDM Packet-switched network architecture with support for group communication," To appear in proc. of IEEE International Conference on Networking (ICN01), Colmar, France, July 10-13, 2001.

/5/ Shrikhande, et al., "HORNET: Packet-over-WDM Multiple Access Metropolitan Area Ring Network," IEEE J. Selected Areas Commun. 18, 11, (Oct. 2000), pp. 2004-2016.

/6/ J. Anderson, et al. "Protocols and Architectures for IP Optical Networking," Bell Labs Technical Journal, Jan. - Mar. 1999, pp. 105-124.