OWI6.pdf

Optical Flow Switching[1]

Vincent W.S. Chan
Joan and Irwin Jacobs Professor
Claude E. Shannon Communication and Network Group, Research Laboratory of Electronics
Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology
Email: chan@mit.edu

Abstract. *Present-day networks are being challenged by dramatic increases in bandwidth demand of emerging applications. We will explore a new transport, "optical flow switching", that will enable significant growth, power-efficiency and cost-effective scalability of next-generation networks.*

Summary. In the early days of the Internet, the most precious resource was long-haul transmission capacity. The electronic packet switching (EPS) architecture was designed to use this resource as efficiently as possible Figure 1. The problem with the EPS architecture, is its scalability: to keep apace with the unfolding exponential growth in data network bandwidth demand. Even electronic processing advancing with Moore's Law will not be able to avert a crisis in electronic-based switching.

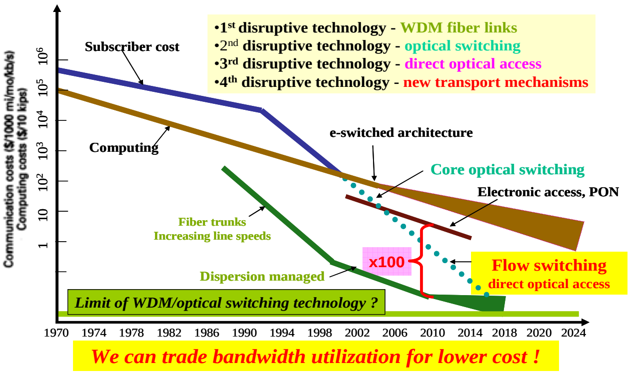

Figure 1. Network cost reductions owing to disruptive optical technologies, [1].

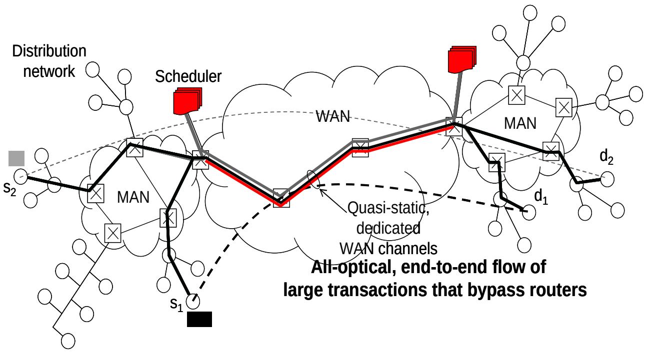

Figure 2. OFS with transparent, end-to-end data flow between users.

In this paper, we explored the use of an optical network transport mechanism, Optical Flow Switching (OFS), for lowering network cost by the same order of magnitude as anticipated future increases in data volume per transaction (two to three orders of magnitude). OFS is a scheduled, all-optical, end-to-end service in which connections are established in response to flow-based requests by client-layer IP network schedulers (e.g., routers) for direct access by individual users, Figure 2. To use network resources efficiently, service holding times of wavelength channels are required to be on the order of hundreds of milliseconds or longer. Scheduling of flows with a time horizon of several transactions also help to achieve high network utilization albeit with some queueing delay at the entrance of the network usually in the form of holding the user from transmission until the network is ready. However, there are specialized applications that have stringent time deadline and would not like to wait in a queue but will be willing to pay for more expense to use the network with short notice such as a single roundtrip time from source to destination. In this paper we will focus on the major factors on the network architecture that result in substantially lowering the cost of providing the OFS transport mechanism. We will also explore a fast OFS for stringent time deadline services. This fast flow switching architecture is highly efficient and economical for on-demand high data rates transfers, distributed sensor data ingestion, as well as distributed computing and processing with short time deadlines and bursty, high-volume data transactions. There is a tradeoff among three observable network performance parameters: delay, blocking probability, and wavelength utilization. The key to high utilization of backbone wavelength channels – a precious network resource owing to the necessity of optical amplifiers and dispersion management – is statistical multiplexing of large flows from many users in a scheduled fashion. Thus, high network utilization can be achieved if the users are willing to wait for service according to a schedule (incurring delay) or accept high blocking probability upon request for service. Figure 3 shows the increase in utilization for the same blocking performance as the scheduling horizon increases.

There are 5 major features that contribute to the OFS architecture that lowers its cost by ~ 50 times versus IP packet switching for large transactions:

1. *End-to-end flows – reducing access and core router traffic by 99%.*

---

[1] This work has been supported by NSF, DARPA and CISCO

2. *The use of optimized logical MAN topology – minimizing active component (especially optical switch) costs.*
3. *Remotely pumped optical access LAN to increase the number of access users/LAN in order to share expensive head-end equipment more cost efficiently.*
4. *Efficient MAC protocol for fast time statistical multiplexing without hgih speed switching and contention collisions.*
5. *The use of quasi-static network management and control in the core (WAN) for greatly simplified and scalable network control plane.*

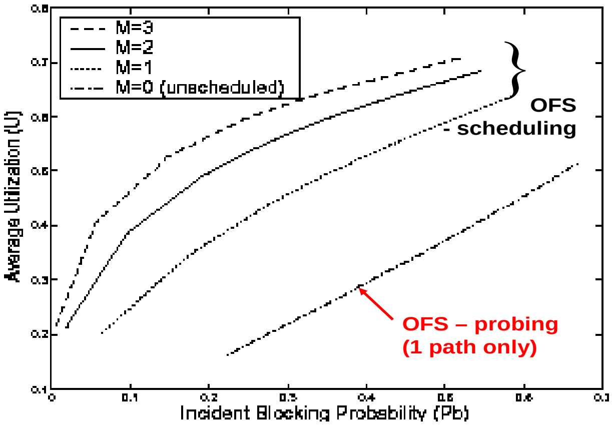

Figure 3. OFS with and without scheduling horizons, [19], showing increased utilization with increasing scheduling horizon (M transactions) for the same blocking probability performance.

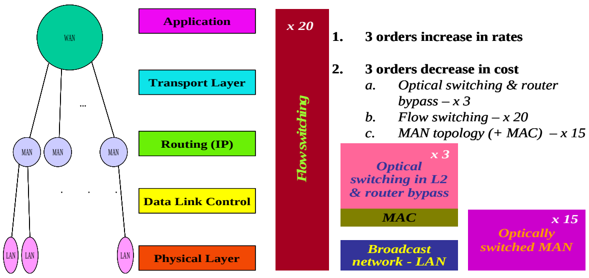

End-to-end flows	*Router/e-switch bypass, reduces router traffic by 99%*
Quasi-static WAN	*NM&C scalability – distributed fast scheduling*
Opt MAN mesh	*Minimizes expensive active components*
Pumped access net	*Increases statistical sharing for efficiency*
MAC scheduling	*Fast time statistical sharing w/o switching & collision*

Figure 4. 5 major architecture features that allow OFS offer ~50 times lower cost to large transactions compared to IP packet switched service.

The results of the throughput-cost comparison are given in Figure 5, [1,2,4,5,13,15]. OFS is the most scalable architecture of all when the average user data rate is high and the number of users in the network is large; EPS is most sensible at low to moderate data rates with a small number of users; the GMPLS architecture, which is conceptually intermediate to EPS and OFS, is optimal at moderate user data rates with a moderate to large number of users; and, finally, there does not exist a regime of optimality for Tell-and-Go, TaG, since the low cost of scheduling in OFS yields great performance benefit relative to the otherwise identical TaG architecture.

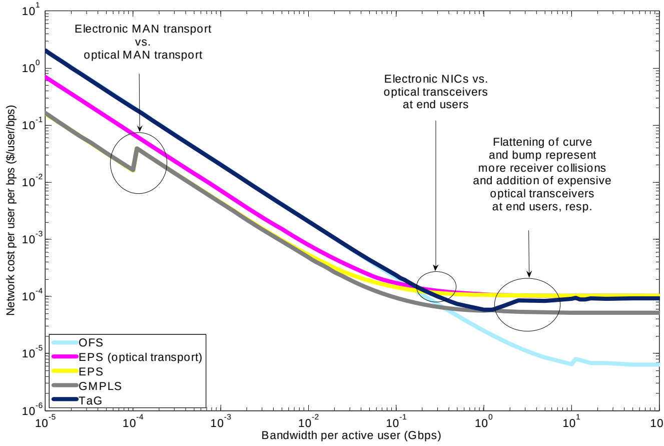

Figure 5. Cost/user/bit as a function of user rates for different transport mechanisms, [1].

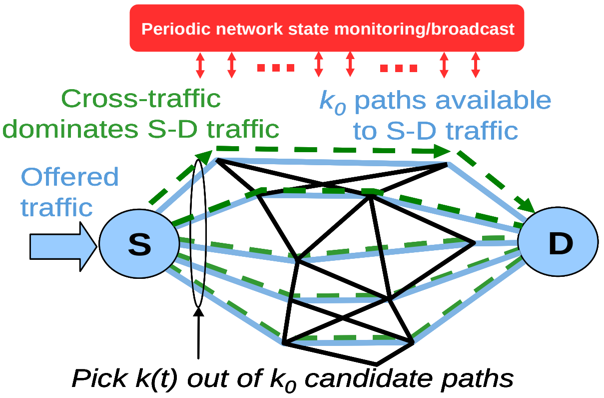

Figure 6. The (slow) centralized network management system periodically broadcasts the path states at regular intervals. The user/scheduler probes $K_p(t) = \min\{K_k, K_o(t)\}$ paths to achieve the desired blocking probability.

OFS can serve high performance applications with very stringent time deadline constraints; setup times only need be slightly longer than one roundtrip time with networks at moderate to high loading. Figure 6 shows the performance of a MAC protocol that a multi-path probing mechanism using periodically updated network state information and with pre-emption of lower priority traffic. The updating scheme calls for a slow control plane, which refreshes and broadcast network states only periodically on the order of seconds or longer. This algorithm, a combination of both slow centralized and fast distributed processes, delivers an efficient and scalable control design, providing different classes of OFS services based on desired delay/cost trade-offs and will allow OFS achieve high utlization > 80% with good quality of service. Thus, the proposed transport mechanism, OFS, with its new network architecture features, that exploit the strengths of optics to serve large transactions, will enable orders of magnitude of cost reductions. The consequence of a shift towards OFS is that most architectural elements of the network – from the physical layer to the higher network layers, as well as network management and control, must be redesigned at the fundamental level.

Background References

[1] Guy Weichenberg and Vincent W. S. Chan, “Design and Analysis of Optically Flow Switched Networks,” IEEE/OSA Journal on Optical Communications and Networking, Aug. 2009.

[2] Kyle C. Guan and Vincent W. S. Chan, “Cost-Efficient Fiber Connection Topology Design for Metropolitan Area WDM Networks,” IEEE/OSA, Journal on Optical Communications and Networking, June 2009.

[3] G. Weichenberg, V.W.S. Chan, and M. Medard, “On the Capacity of Optical Networks: A Framework for Comparing Different Transport Architectures,” OCNIEEE Journal on Selected Areas in Communications, Volume 25, Issue 6, pp. 84 – 101, August 2007

[4] V.W.S. Chan, “Editorial: Optical Network Architecture from the Point of View of the End User and Cost,” IEEE Journal on Selected Areas in Communications, Optical Communications and Networking, Volume 24, Issue 12, pp. 1-2, December 2006.

[5] C. Guan and V.W.S. Chan, “Topology Design of OXC-switched WDM Mesh Networks,” IEEE Journal on Selected Areas in Communications, Optical Communications and Networking, Volume 23, Issue 8, pp. 1670 – 1686, Aug. 2005.

[6] V.W.S. Chan, S. Chan, S. Mookherjea, "Optical Distribution Networks," Optical Networks Magazine, Volume 3, Number 1, pp. 25-33, January/February 2002.

[7] M. Kuznetzov, N.M, Froberg, S.R. Henion, H.G. Rao, J. Korn, K.A. Rauschenbach, E.H. Modiano, V.W.S. Chan, “A Next-Generation Optical Regional Access Network,” IEEE Communications Magazine (Invited), Volume 38, Issue 1, Pages: 66-72, Jan. 2000.

[8] V.W.S. Chan, K.L. Hall, E. Modiano, K.A. Rauschenbach, “Architectures and Technologies for High-Speed Optical Data Networks,” IEEE Journal of Lightwave Technology (Invited), Volume 16, Issue 12, pp. 2146 – 2168, Dec. 1998.

[9] I.P. Kaminow, C.R. Doerr, C. Dragone, T. Koch, U. Koren, A.A.M. Saleh, A.J. Kirby, C.M.Ozveren, B. Schofield, R.E. Thomas, R.A. Barry, D.M. Castagnozzi, V.W.S. Chan, B.R. Hemenway, Jr., D. Marquis, S.A. Parikh, M.L. Stevens, E.A. Swanson, S.G. Finn, R.G. Gallager, "A Wideband All-Optical WDM Network," IEEE Journal on Selected Areas in Communications, Vol. 14, Issue 5, June 1996.

[10] R.A. Barry, V.W.S. Chan, K.L. Hall, E.S. Kintzer, J.D. Moores, K.A. Rauschenbach, E.A. Swanson, L.E. Adams, C.R. Doerr, S.G. Finn, H.A. Haus, E.P. Ippen, W.S. Wong, M. Haner, "All-Optical Network Consortium - Ultrafast TDM Networks," IEEE Journal on Selected Areas in Communications, June 1996.

[11] S.B. Alexander, R.S. Bondurant, D. Byrne, V.W.S. Chan, S. Finn, R. Gallager, B. Glance, H.A. Haus, P. Humblet, R. Jain, I.P. Kaminow, M. Karol, R.S. Kennedy, A. Kirby, H.Q. Le, A.A.M. Saleh, B. Schofield, J. Shapiro, N.K. Shankaranarayanan, R.W. Thomas, R. Williamson and R. W. Wilson, "A Precompetitive Consortium on Wideband All-Optical Networks," IEEE Journal of Lightwave Technology, May 1993.

[12] Guy Weichenberg, Vincent W. S. Chan and Muriel Medard, “Performance Analysis of Optical Flow Switching”, ICC Dresden Germany June 2009.

[13] Guy Weichenberg, Vincent W. S. Chan, Eric A. Swanson, Muriel Médard, “Throughput-Cost Analysis of Optical Flow Switching”, Optical Fiber Communication Conference, San Diego, CA, April, 2009.

[14] G. Weichenberg and V.W.S. Chan, “Access Network Design for Optical Flow Switching,” IEEE Global Telecommunications Conference, 2007, Washington, D.C., Nov. 2007

[15] V.W.S. Chan, G. Weichenberg and M. Medard, “High-Impact Research Problems in Optical Networking” 16th International Conference on Computer Communications and Networks (ICCCN), August 2007.

[16] G. Weichenberg, V. W. S. Chan, and M. Medard, "On the throughput-cost tradeoff of multi-tiered optical network architectures," IEEE Global Telecommunications Conference 2006, San Francisco, Nov. 2006.

[17] G. Weichenberg, V. W. S. Chan, and M. Medard, "Cost-Efficient Optical Network Architectures," 32nd European Conference on Optical Switching (ECOC 2006), Cannes, France, September 2006.

[18] G. Weichenberg, V. W. S. Chan, and M. Medard, "On the Capacity of Optical Networks: A Framework for Comparing Different Transport Architectures," IEEE Infocom, Barcelona, April 2006.

[19] B. Ganguly, V.W.S. Chan, “A scheduled approach to optical flow switching in the ONRAMP optical access network testbed,” Optical Fiber Communication Conference 2002, Mar. 2002.

[20] S. Mookherjea, V.W.S. Chan, “Remotely-Pumped Optical Distribution Networks,” Conference on Lasers and Electro-Optics 2001, May 2001.

[21] V.W.S. Chan, A.J. Kirby, A. Saleh, "AT&T/DEC/MIT Consortium on Wideband All-Optical Networks," LEOS 1993, San Jose, CA, November 1993.

[22] V.W.S. Chan, "All-Optical Networks," IEEE 1993 Optical Fiber Communications Conference, San Jose, California, February 1993.

[23] V.W.S. Chan, "All-Optical Networks," LEOS Summer Topical Meeting (Invited), Santa Barbara, California, August 1992.