Optical Switching and Networking 6 (2009) 37–43

Contents lists available at ScienceDirect

Optical Switching and Networking

journal homepage: www.elsevier.com/locate/osn

Hierarchical time-sliced optical burst switching ☆

Vijay Sivaraman, Arun Vishwanath *

School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, NSW 2052, Australia

ARTICLE INFO

Article history:
Received 19 March 2008
Accepted 17 May 2008
Available online 18 June 2008

Keywords:
Optical networks
Optical burst switching

ABSTRACT

To overcome the need for large buffers to store contending bursts in optical burst switched (OBS) networks, a recent variant called time-sliced OBS (TSOBS) suggested that bursts be sliced and spread across multiple frames of fixed-length time-slots. Since TSOBS is rigid in its frame structure, this paper generalises TSOBS to allow a hierarchy of frames. Termed hierarchical TSOBS (HiTSOBS), this scheme supports several granularities of rates, and permits multiple traffic classes with different loss-delay requirements to efficiently share the network. Our contributions are as follows: first, we present an architecture for HiTSOBS and offer it as a viable option for the realisation of flexible and cost-effective OBS networks. Second, we develop mathematical analysis to study the loss and delay performance of the proposed HiTSOBS system. Finally, we present simulation results that captures these loss-delay tradeoff values. Our HiTSOBS architecture gives network operators the freedom to choose the right mix of traffic with desired loss-delay requirements to coexist in the network.

1. Introduction

Wavelength Division Multiplexing (WDM) optical networks provide enormous bandwidth and are promising candidates for information transmission in next-generation high-speed networks. It is possible to realise 10–40 Gbps bandwidth on a single wavelength in commercial WDM networks today. However, a fundamental concern in the continued scalability of optical networks is the huge disparity in switching speeds between optical and electronic switches in the core of the network. With a vision towards evolving to an all-optical Internet, optical switching can be classified into three categories — optical circuit switching (OCS), optical packet switching (OPS) and optical burst switching (OBS).

In OCS networks, lightpaths are used to transmit data between two end nodes [1,2], where a lightpath is defined as an all-optical circuit switched medium with possible wavelength conversion at the intermediate nodes along the transmission path. Although OCS is easy to implement, it suffers from poor statistical multiplexing gains if the source node does not have any data to send, thereby leading to poor resource and bandwidth utilisation.

OPS [3,4] on the other hand is similar to traditional electronic packet switching, wherein packets are switched directly in the optical domain, without the need for any electronic conversion. However, the most important concern is contention, which occurs at a switching node whenever two or more packets try to leave on the same output interface, on the same wavelength, at the same time. Unlike in electronic RAMs where as many as a million packets can be buffered during times of contention, buffering in the optical domain remains a very complex and expensive operation. Spools of fibre can implement fibre delay lines (FDL) that can buffer light by delaying the signal, however the size of the optical crossbar increases with bigger FDLs, thereby making all-optical switches very expensive. Recent research work [5–8] explores the feasibility and performance of transport protocols for realising OPS networks in routers equipped with very small buffers, i.e., only a few dozen packet buffers that can be

☆ A short summarized version of this paper was presented at the First Symposium on Advanced Networks and Telecom Systems (ANTS) 2007 in Mumbai, India, in December 2007.

* Corresponding author.
E-mail addresses: vijay@unsw.edu.au (V. Sivaraman), arunv@ee.unsw.edu.au (A. Vishwanath).

1573-4277/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.osn.2008.05.005

implemented in an on-chip optical memory. This remains an active area of research, and if successful, could lead to commercial large-scale deployment of OPS networks in the future.

Optical Burst Switching (OBS) [9,10] is a hybrid of circuit and packet switching: aggregates of packets, called bursts, are switched atomically within the network, while a control packet is sent ahead of the burst to set up a short-lived end-to-end circuit for the burst. OBS thus combines the scalability of optics for fast data plane switching with the flexibility of electronics for switching decision control. An unfortunate consequence of this architecture, similar to OPS, is that the optical buffering required for contention resolution grows in proportion to the burst size. The control-plane advantage of large bursts is thus tempered by the larger buffers required in the data-plane.

1.1. Time sliced optical burst switching

A variant of optical burst switching, called Time Sliced OBS (TSOBS), was proposed in [11] to overcome this problem of having larger buffers. Time is divided into frames that contain a given number of fixed-length slots. TSOBS slices a burst, and transports successive slices in the same slot location of successive frames. This preserves the control-plane scalability of OBS (since only one switching decision is required to switch all slices belonging to a burst), while drastically reducing the optical buffering required at switching nodes (since a contending burst need only be buffered a slice at a time, independent of burst size). In addition, since switching is entirely done in the time domain rather than the wavelength domain, TSOBS eliminates the need for having wavelength converters, which substantially reduces the cost of designing such a network. Further, the authors identify three important factors that affect the cost and performance of optical time-slot interchangers (OTSI), which is a key component of the TSOBS system. They are the size of its internal crossbar, amount of fibre needed for the FDLs to reorder the timeslots and the number of switching operations that a burst may be subjected to when passing through the OTSI. Several blocking and non-blocking architectures for implementing the OTSIs are also proposed and analysed.

1.2. Related work

Following the TSOBS system, the work in [12,13] proposes a variant called Time-Synchronized Optical Burst Switching (SynOBS), which not only assumes the presence of fibre delay lines, but also considers the impact of full wavelength conversion. Several FDL reservation mechanisms - core node without FDLs, separated, shared and multi-length FDLs, are proposed and analysed using discrete time Markov chains to compute the burst drop probability. Their study also suggests that timeslot size must be chosen with care to achieve the best timeslot utilisation, which subsequently reduces burst blocking probability. In [14], the authors consider a slotted optical burst switching network (SOBS) and argue that such a network improves the overall link utilisation. They claim that their work is the first to point out the advantage of SOBS in supporting Quality of Service (QoS) requirements, and also propose a new cost-effective method for aligning packets at core nodes.

Akin to TSOBS but termed all-optical cell switching, [15] proposes FDL assignment algorithms to achieve low cell-loss rate to support both guaranteed and best-effort traffic. In [16], an analytical model is developed to estimate the overall blocking probability for a multi-fibre TSOBS network. The model is able to compute the overall blocking probability for circuit switched, best effort and multi-class traffic services in the network. Their results indicate that multi-fibre TSOBS can achieve the same level of performance (with respect to blocking probability) as a conventional OBS network (employing just-in-time reservation protocol [17]) with wavelength conversion functionality.

To address the fairness issue in OBS networks, [18] presents a new scheduling algorithm using round-robin scheduling, termed Almost Strictly Proportional Fair Scheduling (ASPFS), for SOBS networks with full wavelength conversion capability. SOBS is chosen to overcome the difficulty of the lack of large optical buffering in today's optical networks. Analytical and simulation results indicate that ASPFS is a promising candidate to provide fairness in future OBS networks. Slot allocation for TSOBS networks using centralised control is discussed in [19]. Request to calculate a path and an appropriate slot for a burst from an ingress OXC (Optical Cross Connect) is delivered to a centralised controller, which then computes these values. At the expense of an increased queueing delay at the ingress node, their scheme is able to improve channel utilisation, which is derived using both analysis and simulation. In [20], the authors propose a scheme to balance the loss-delay tradeoff in a slotted optical packet switched network. Using analysis and experimental results, the authors study the effect of ingress traffic conditioning, i.e., the effect of spacing out optical packets that feed into an OPS core node. They demonstrate that such a scheme can effectively bring down the packet loss probability to acceptable levels even when only minimal buffering is available at the core node. However, this low loss comes at the cost of an increased end-to-end delay of the conditioned traffic flow. The resulting strategy allows network service providers to choose the appropriate loss-delay values for operating their networks.

1.3. Our contributions

While TSOBS successfully addresses the scalability of optical burst switching systems, it is excessively rigid in its frame structure. The frame size (i.e. number of slots per frame) is a key parameter that has to be universally pre-configured at all switches. A small frame size increases contention probability since overlapping bursts are more likely to pick the same slot number, while large frame sizes induce larger end-to-end delays due to each flow having access to a reduced fraction of the link capacity (one slot per frame), leading to significant queueing delay at the ingress edge node. This loss-delay trade-off, determined by frame size, is uniform across all traffic flows, and cannot