OPTICAL BURST SWITCHING FOR NEXT GENERATION TRANSPORT NETWORKS

MASTERS-TO-PHD CONVERSION REPORT

JOLYON WHITE

SUPERVISORS: PROF. ROD TUCKER AND DR. KEPING LONG

ARC SPECIAL RESEARCH CENTRE FOR
ULTRA BROADBAND INFORMATION NETWORKS
THE DEPARTMENT OF ELECTRICAL AND ELECTRONIC ENGINEERING
THE UNIVERSITY OF MELBOURNE

13 MARCH 2002

Contents

1. Introduction

For a number of years now it has been well recognised that the Internet has been growing at exponential rates, doubling in traffic volume every 9 months before 1997 and every 6 months since 1997 [1]. This traffic doubling has both motivated and been driven by advances in technologies. For instance in the early 1990s, emerging software technologies, primarily hypertext and its supporting protocols which make up the World Wide Web, made the Internet useful and accessible to the broad public. At the same time, the introduction of optical fibre as a transmission medium, followed by the use of Wavelength Division Multiplexing (WDM), began a rapid decrease in the cost of transmission bandwidth which has made the Internet highly affordable to ordinary people. The affordability and raw computing power of home PCs have made sophisticated applications (such as real-time video and audio) available to many people. Accessing content for such applications over the Internet has been identified as a key potential revenue source for Internet Service Providers (ISPs) and carriers, which has fuelled the drive towards providing high bandwidth commodity access services to customers. Carriers also see further gains in delivering many disparate services over the one core network, combining voice, data, video, telepresence, and future services through a single infrastructure.

If this predicted explosion of demand for high bandwidth access eventuates, carriers will be required to increase the bandwidth of their core backbone networks by many orders of magnitude to account for the massive aggregate traffic volume envisaged, a fact compounded by the empirically observed bursty nature of Internet traffic [2]. Optical fibre and photonic technologies are the prime candidates for providing these capacity increases. Optical fibre offers enormous bandwidth in the telecommunications wavelength band around 1550nm, estimated at 25Tbits/sec [3]; they are far more compact than copper cables, with many fibres able to fit in a single underground duct; they provide very low loss, allowing long repeater-less spans, which reduces maintenance costs. They also allow multiplexing of many signals at different centre wavelengths onto a single fibre (WDM). Fine wavelength spacing of these signals is known as Dense Wavelength Division Multiplexing (DWDM) which is expected to be deployed widely in the future and will produce significant increases in available bandwidth to meet demand. For example, whereas coarse WDM might feature 8, 16 or 32 wavelengths on a single fibre, DWDM systems are expected to feature more than 100 wavelengths per fibre.

2. Problem Statement

Currently deployed networks consist of point-to-point (D)WDM links. This means that at each switching point, the optical signal is converted to electrical form, and the processing and forwarding is done in the electrical domain. This is known as an Optical-Electronic (O/E) conversion. When the signal is passed to the outgoing port, it is again converted and modulated onto the fibre as an optical signal (Electronic-Optical or E/O). Such a switching point is said to perform O/E/O conversion (optical-electronic-optical).

In such a network, all communication is limited by the capabilities of the electronics in the system. Once the optical signal is converted into an electronic signal, some form of processing unit must examine the data and decide to which output fibre and wavelength the data must be forwarded. The data must then be remodulated onto the optical carrier. All of these steps involve significant delays enforced by the limitations of electronic processing speed. It is complicated by the fact that the full data payload of the incoming signal must be read and stored before being forwarded. In a packet switch (or router) the burden of the electronics is increased because the forwarding decision must be made for each packet at each hop and it is an expensive operation. Variable length packets must also be buffered at the router, requiring large amounts of electronic memory and sophisticated buffer management algorithms.

A question comes to mind at this point: *Can we achieve high switching throughput if an optical switching technology is used?* In other words, can we use optical technologies in the switch fabric, instead of electronic technologies? There currently exist numerous candidate technologies which could be used to switch data in the optical domain without need for converting the payload data into the electronic domain. They have their own advantages and limitations, some of which are significantly different to those found in the electronic domain. Principally, the state of the art in optical processing is currently embryonic, and so complicated control algorithms must be implemented in electronics still – thus at least the control information (e.g. packet headers) must still be converted to electronic form for processing. Secondly, optical buffering technologies do not yet exist beyond simple delays or recirculating loops, which poses problems when contention for limited outgoing transmission resources is experienced.

Given that we have available a number of optical switching technologies, how can we use the advantages of these technologies and overcome their drawbacks to achieve high throughput optical switches? How can we build a transport network out of such switches to provide service guarantees to the diverse applications expected in the next generation networks?

This research aims to find some answers to these questions. The specific focus of this research is on Optical Burst Switching (OBS) as a compromise between the demands of forecast traffic (which mandate the statistical sharing of packet switching) and the realities of optical technologies (for which circuit switching is much easier to implement). OBS is a possible stepping stone on the way to Optical Packet Switching (OPS). Exploring the continuum of switching techniques that exist between full circuit switching and true OPS is another goal of this research, which it is hoped will provide insight into the best path to take in order to advance towards an all-optical packet switching network.

3. Literature Review

This section provides a detailed overview of the literature pertaining to Optical Burst Switching. In so doing, it serves to provide motivation for OBS as an alternative to existing optical switching techniques and to place

OBS in its proper context among the current research avenues in optical networks.

3.1. **All Optical Networks.**

Optical networks can be divided between first generation and second generation networks. First generation networks were point-to-point networks which used fibre as a faster substitute for copper cable. Wavelength Division Multiplexing technology extended these systems to have more than one wavelength per fibre, increasing their capacity several times over. At each switching point in such networks, all wavelengths are terminated and converted to electrical form, and are remodulated onto the optical carrier at the output [4]. Second generation networks obviate the need for conversion to the electronic domain by providing switching and routing services at the optical layer. These networks are known as *all optical networks*.

Similarly, all optical WDM networks can be classified along three lines: these are the *broadcast and select* architecture, the *wavelength routing architecture* and the *photonic* or *optical packet switching architecture* [4]. Whilst broadcast-and-select networks (also known as *single-hop networks* [5]), which are based on passive optical components, find their natural place in the access networks arena, wavelength routing and optical packet switching networks are intended for the metro and core transport networks.

An overview of the development of such networks, including broadcast-and-select and wavelength routing networks, can be found in Mukherjee [6]. Mukherjee also discusses such issues as network control, fault management and protection, multicasting and traffic grooming in all-optical networks. Okamoto gives an overview of all optical networks from an architectural and management viewpoint in [7]. Broad design and performance issues encountered in WDM transport networks are addressed by Karasan and Ayanoglu [8].

3.2. **Wavelength Routing Networks.**

Optical networks which perform some optical routing and involve traversal of multiple switches for typical connections are called *wavelength routing networks* [4]. As stated in [4] (p. 507), one of the most important drivers behind wavelength routing is the desire to reduce the processing overhead of pass-through traffic. In many applications, for example SONET or SDH networks, much of the electronic processing at an add-drop multiplexer (ADM) is dedicated to forwarding pass-through traffic which is not destined for the node in question. As more wavelengths are added to a WDM point-to-point network, each wavelength requires its own terminating equipment, and so the problem is compounded. If a full wavelength of capacity is dedicated to pass-through traffic, then it makes sense that it should be possible to forward it in the optical domain, allowing the elimination of a full SONET/SDH ADM and the significant costs with which it is associated.

Chlamtac et. al. [9] introduced the concept of a *lightpath* from a source node to a destination node. Under this idea a node which needs to send data requests a lightpath to the destination node, which is set up by signalling all of the optical switches along the path with the request. If a switch accepts the request then it dedicates a path through its switch fabric to carry the data for the request between the incoming and outgoing fibres, and from then on it knows that it can identify that connection by the wavelength of the incoming signal on the input fibre. This means that there is no need for the switch to perform complex processing of frame headers. Furthermore, wavelength selective technologies can be used to perform the switching function, and to avoid electronic conversion and buffering of the payload. The lightpath concept has since been studied extensively. Key analyses of blocking processes in lightpath-based networks can be found in Birman [10], Barry and Humblet [11] and Kovačević and Acampora [12]. The problem of assigning a wavelength to a lightpath along its route is a key problem in wavelength routing networks, known as the Routing and Wavelength Assignment problem. Various solutions have been proposed, see [8] and [13]. Birman points out that wavelength routed networks are not equivalent to electronic circuit switched networks because circuits in a circuit switching network are indistinguishable whereas wavelengths are not because of the continuity constraint that a lightpath take the same wavelength on all links. For this reason, wavelength routed networks suffer higher blocking probabilities than circuit switched networks.

Yates et al [14] studied networks where wavelength conversion is provided by exploiting Four Wave Mixing (FWM) in semiconductor optical amplifiers (SOAs). Such wavelength converters are *limited*, meaning that they degrade the signal by an amount which depends on the difference between input and output wavelengths. Thus they do not have absolute freedom to convert from any input to any output wavelength. Yates et al introduced the notion of a cost incurred in traversing a cross-connect and a budget for lightpaths from any given point to a given destination. Using this framework they showed that limited wavelength conversion networks can give blocking performance close to that of networks with ideal converters.

3.3. **Optical Packet Switching.**

Wavelength routing networks are designed to provide coarse-grained capacity on demand which is intended to be established and maintained for a relatively long time. This is because wavelength routing is a form of circuit switching, and in a circuit-switched network an end-to-end connection must be established between source and destination hosts before data may be transferred [15]. The forwarding decision is thus made at the call setup time and during the lifetime of the call the switch must read, store and forward each frame of received data for that call. Thus the overhead is large at the time of call setup (which requires a round-trip-time) and smaller during the call life-time. However, calls are intended to be long compared to setup and teardown times so that efficiency is high if the bandwidth is fully used during the life of the call. This is good for optical networks because optical switching technologies are generally slow compared to electronic technologies. Fast optical technologies do exist but they have poor loss and crosstalk characterists which limit their use.

In general it can be difficult to utilise a circuit efficiently, due to such effects as daily traffic fluctuations seen by most ISPs, and the need to switch traffic in finer grained quantities as occurs in metro networks [16]. Moreover, with Internet Protocol (IP) emerging as the dominant layer 3 protocol, and carriers moving towards deploying extensive IP networks, it makes sense to ask whether wavelength routing is a sensible long-term proposition for the underlying transport network. As pointed out by Tucker and Zhong [17], one should ask whether packet switching services should be delivered by the transport network itself, in this case an optical transport network.

In such an optical (or photonic) packet switching network, packets are switched in all-optical form. Tucker and Zhong [17] provide an extensive overview of photonic packet switching. They identify the functions of a photonic packet switch as packet synchronisation, packet routing, packet header replacement and packet buffering. Packets in such networks generally have fixed duration and the network is slotted, requiring proper synchronisation of packets at the switch inputs and recovery of timing information. These functions must be provided photonically, transparent to bit rate, modulation format and higher-layer protocol encoding.

Spring et al [18] demonstrated a photonic header replacement technique which avoids a dispersion-induced timing error which would be caused if a laser was used to generate the new header at the switch, since the wavelength of the incoming packet and replacement header would not be exactly the same. They used continuous wave light transmitted with the packet together with clever use of delays and optical gates to modulate the new header onto the original laser source light.

The recirculating buffer is one of several possible optical packet buffer architectures. Another architecture consists of a number of *fibre delay lines* (FDLs) of varying lengths. Each length is usually a multiple of the packet size [17]. The FDLs may be arranged in parallel or in series, allowing a large number of flexible delays to be selected by routing the packet through the correct path in the switch [19, 20]. Recirculating buffers suffer from the fact that at each recirculation the attenuation of the packet in the fibre must be reversed by amplification, which introduces Amplifier Spontaneous Emission (ASE) noise. Travelling type buffers are inflexible since only a finite number of delay times can be synthesised with a finite number of FDLs. Hence photonic buffering remains a significant problem which lags behind the logical and temporal capabilities of electronic memories.

The question of which among wavelength routing or photonic packet switching is the best long term technology has not yet been answered fully. However, Roberts [1], one of the designers of the original Internet, provides some insight. When the Internet was being developed, he contends, communications link cost was falling slowly (halving every 79 months) compared to computer processing equipment costs (halving every 21 months). The two costs "crossed over" in 1969, the year the first Internet hosts were deployed, such that it made more sense for computing power to be concentrated at the nodes to provide efficient statistically multiplexed packet switching rather than spend money on higher communications link bandwidth which would be inefficiently used by the less complex circuit switching.

These trends continued, he says, until 1997 when link costs began halving every 12 months when WDM systems were first rolled out. The question remains as to whether links will again become cheaper than processig power, making circuit switching more attractive. Roberts speculates that the full potential of WDM technology will be realised by 2008, at which point it will be increasingly more difficult to cram more wavelengths into a fibre, and link costs will stabilise at at least an order of magnitude more than processing costs. He argues that this gives the ascendency to packet switching. Whether or not the future will eventuate in this manner is not clear, however Roberts does at least provide criteria for making the decision if it becomes necessary.

3.4. **Optical Burst Switching.**

A new switching scheme called optical burst switching (OBS) has been proposed to provide flexible bandwidth packet-like services whilst overcoming some of the limitations faced by both optical packet switching and wavelength routing. OBS is a hybrid switching scheme that resembles circuit switching and wavelength routing in some respects, and packet switching in others. OBS was first introduced by Yoo and Qiao in [21, 22], and builds on the Tell-And-Go protocol (TAG) developed by Hudek and Muder [23]. These protocols are based on work done in fast circuit switching for Asynchronous Transfer Mode (ATM) networks where a virtual circuit is set up on-the-fly for a burst of cells which cut through intervening switches without buffering, and without waiting for acknowledgement of the allocation of the circuit. Thus if the circuit is not successfully set up, the burst is lost at the switch where the setup failed This was not a successful approach in ATM because of the electronic complexity required, however with optical switching technologies much of this complexity is avoided.

In these protocols, the network is assumed to be a WDM network with a number of wavelengths carried on each fibre, and optical switching technologies utilised in the switching nodes. In burst switching, a number of fundamental protocol data units (e.g. IP packets or ATM cells) are grouped into a burst for transimission as a single indivisible block. The idea here is to amortise the switching overhead and protocol processing overhead over a larger amount of payload data, thus enabling cheaper and less capable switches to be employed.

In the TAG protocol of Hudek and Muder, a burst of data is transmitted preceded by a control packet which attempts to reserve bandwidth and configure switches, and is followed by a control message to tear down the path. There is no acknowledgment of the path setup, and the control messages are sent on a separate channel. In the *just enough time* (JET) protocol of Yoo and Qiao [21], the control packet precedes the burst by an offset time T, which gives the switch a "look-ahead" capability and allows it to intelligently schedule the burst of data onto a wavelength on the desired output fibre. Furthermore, the capacity is reserved from the time when the burst data reaches the switch, not when the control packet arrives. This means the intervening link time can be allocated to some other transmission, improving utilisation. This is known as "delayed reservation". Also, the control packet contains the duration of the data burst so that the switch knows when the capacity will be free again, and know teardown message is required. Hence JET has reduced the setup and teardown costs compared to circuit switching and TAG, and is therefore much better suited to statistically multiplexed short timescale switching.

Yoo and Qiao found that their protocol performed significantly better than TAG in terms of utilisation, however in [22] they presented results which showed that OBS has very high burst dropping probabilities which will induce significant unwanted retransmissions and bandwidth wastage. They present a scheme in [22, 24] which differentiates between classes of service based on the offset time T assigned to each class. The lowest priority class gets a base offset $T_0 = T$, and every higher class i has an offset time $T_i = T + \sum_{j=1}^{i} \Delta T_j$, for $i \geq 1$ and $\Delta T_i > 0$ for all $i \geq 1$. Higher priority classes thus see further into the future and thus have more chance of finding a wavelength free for the duration of their bursts. In [22] they showed that their scheme reduces the burst dropping probability by several orders of magnitude for high priority bursts at the expense of reducing the blocking probability slightly for low priority bursts, in a two class system. They also found that the overall blocking probability is unchanged when compared to a classless system. They extended their work to analyse multiple classes and systems with FDL buffers in [24] where they applied $M/M/k/k$ and $M/M/k/D$ models to provide upper and lower bounds on the burst blocking probability for each class. They found that significant isolation can be achieved for the highest priority class, at the expense of the lower classes. This scheme obviously introduces extra delays for high priority classes which may be at odds with low delay requirements of some services. A good overview of Yoo and Qiao's work is found in [25]. A detailed discussion of the motivation for the JET protocol, together with a simulation-based performance evaluation for networks which retransmit lost bursts can be found in [26].

In parallel with Qiao and Yoo, Turner [27] developed a similar burst switching technique which is simpler than JET. He calls his technique "Horizon scheduling". This is because the state information in JET consists of the start and end times of each burst already allocated to a wavelength, whereas the Horizon scheme maintains only the current *scheduling horizon* for each channel, i.e. the latest time at which the channel is being used by a burst. Beyond the horizon is where the channel is guaranteed to be free, and bursts may only be allocated there in this scheme. For this reason, it achieves lower utilisation than JET, but it is much simpler to implement and so may be more attractive in eventual commercial implementations.

Wei et al have also devised a scheme for OBS which they call *just in time* (JIT) [28]. In their scheme a wavelength is allocated from the time the control packet is received until the time when a release packet is received, and allocation is granted based on the current state of the data channel, rather than any future state. In this way, JIT is more of a TAG-based scheme and a light-weight approach to the problem which does not achieve utilisations as high as the other two because bursts can be dropped when it might have been possible for them to be accommodated otherwise. Evaulation of their signalling protocol also appears in [29].

Dolzer et al have written an overview of the Horizon, JET and JIT protocols [30], which shows JET outperforming Horizon which outperforms

JIT in terms of blocking probability. They also introduce an analysis of the blocking probability for a 2-class system which follows Qiao and Yoo's quality-of-service (QoS) scheme [24] and extends their analysis. Whilst Qiao and Yoo consider only the high priority traffic load A_H and the low priority traffic load A_L separately when applying Erlang's loss formula, Dolzer et al also consider the portion of carried traffic of the low priority traffic, $Y_L(\delta_H)$, against which the high priority traffic must also compete, since the classes are not perfectly isolated. $Y_L(\delta_H)$ represents the low priority bursts which started transmission prior to the arrival of the high priority control packet and are not completed by the time the high priority burst arrives. This is calculated using the *forward recurrence time* , whose complementary distribution describes the probability that a low priority burst does not finish transmission within the period of length δ_H. This can be calculated from the low priority burst length distribution. Dolzer et al compared their analytical predictions with simulation outcomes and the lower bounds of Qiao and Yoo and showed that their method provides good agreement with simulation, and slightly overestimates the blocking probability.

Whilst JET, Horizon and JIT provide signalling protocols, it remains for each switch to decide to which wavelength an outgoing burst should actually be assigned. This is known as wavelength assignment or wavelength scheduling. Algorithms for wavelength scheduling were not discussed until Xiong et al [31] outlined several algorithms in their paper. In this paper they also described a possible architecture for an optical burst switch, similar in design to existing packet switches, and also proposed a system design for the control packet processor at the switch. They described first a simple first-fit (FF) algorithm which does a round-robin search of availabe wavelengths and assigns the first found; a *latest available unscheduled channel* (LAUC) algorithm which is essentially the same as the Horizon algorithm [27] and assigns the burst to the channel which has the latest horizon time (to use the smallest gaps first so that later arrivals which need larger slots are not blocked); and the *latest available unused channel with void filling* algorithm (LAUC-VF) which exploits knowledge of all burst start and end times to place bursts in gaps between other bursts when it can. Using a self-similar traffic model based on Fractional Gaussian Noise, they simulated their algorithms and found that LAUC-VF performs best, follwed by LAUC and FF, in terms of burst loss ratio.

The LAUC-VF algorithm has since been extended in various ways, for instance Yang et al [32] give a modification to LAUC-VF for hardware implementation and extend the algorithm such that control packets are queued separately according to service class, with the highest priority queue served first (according to the LAUC-VF discipline) and subsequent queues served in descending priority order until all control packets in all queues are exhausted. In this way, the high priority bursts are given the best pick of available wavelengths and are more likely to be accepted, at the expense of the lower priority bursts. This is a simple idea and produces more than an order of magnitude reduction in the blocking probability of high priority bursts (in a 3-class system), again at the expense of a worsening in the blocking probability for the lowest priority class. Their system is also slotted, and the simulation was performed on essentially a ring of core nodes, with each core node (switch) connected to one source as well as its neighbouring core nodes.

A highly novel algorithm is described by Wang et al [33] in which each switch monitors its success rate when sending on each wavelength to each destination, and preferentially transmits on the wavelength which succeeds the most when transmitting to that destination. The switches learn from statistical data collected and use a simple update procedure to update the preference given to each wavelength at each observation step. For low traffic intensity this was shown by simulation to provide up to two orders of magnitude of benefit in blocking probability over random wavelength assignment, however this benefit is all but lost for traffic intensities above 0.25. Furthermore the authors did not compare their algorithm with LAUC or LAUC-VF, which both perform much better than the random wavelength assignment algorithm. Therefore this idea, which is unique, does not appear to perform well enough to compete with other proposed algorithms.

There have been a number of tutorial and overview articles on OBS in various technical magazines recently as well, for instance, [34, 35, 36, 37]. These articles provide architecture overviews and summaries of the findings of the academic journal and conference articles mentioned above, and have raised the profile of optical burst switching. The networks envisaged in these articles consist of a core transport network of OBS switches which forward bursts from their sources to their destinations, which are ingress and egress routers. The ingress routers accept incoming streams of packets and group them into bursts for transmission through the network. For each destination, packets destined for that destination are sent to the same queue for transmission in a single burst. How the ingress router decides when it has enough packets to make up a burst is determined by the *burst assembly process*.

Ge et al first proposed an explicit burst assembly mechanism in [38] in studying the effect of burst assembly as a shaper of the incoming traffic. In their mechanism, for a given destination i, a timer is started when a packet arrives at queue i for destination i and finds queue i empty. When the timer passes a certain amount of time a burst is created and queued for transmission on the data channel. If the burst is too small it is padded with null data. When the burst is transmitted the timer is reset and only starts counting again when another packet arrives for destination i. This algorithm provides an upper bound on delay due to burst assembly and gives bursts a minimum size, albeit at the expense of wasted transmission bandwidth. When they fed their algorithm with bursty traffic generated from a multiplex of a large number of ON/OFF sources with Pareto distributed ON and OFF periods, they showed that the Hurst parameter was reduced from about 0.904 to 0.76 (for a smaller timer threshold) or 0.65 for a larger timer threshold. Hence the traffic launched into the OBS network is significantly less bursty than the input traffic when shaped by the burst assembler.

This indicates that burst assembly can have a very important role to play in the source modelling aspects of OBS and is an area which is open to much further study. Xiong et al proposed a similar burst assembly mechanism which places a limit on both the elapsed time *and* the maximum size of the burst [31].

More recently, a number of papers have appeared which discuss "Wavelength Routed OBS" or WROBS. In this scheme, a wavelength routed optical network is considered in which packets arrive at an edge router (ingress node) and a burst is formed, whereon a request for an end-to-end lightpath is made to the network (centrally managed) and once the lightpath is set up the burst is transmitted (for instance, Düser and Bayvel [39]). Thus the burst must wait for acknowledgement of resources before being transmitted, and this means that this scheme must be subject to large latencies compared to JET based OBS in which the burst does not wait. Furthermore, the burst can still be blocked by the routing and wavelength assignment algorithm, which must be performed very rapidly over the entire network. In their paper, Düser and Bayvel describe a burst assembly algorithm in which a lightpath request is made once a key service quality parameter is exceeded, such as latency or packet loss ratio at the burst assembly buffer. They showed that the arrival and service processes influence the allowable delay at the edge node before packet loss rate (PLR) is exceeded. Bursty processes perform worse than Markovian processes, with edge delays between 28ms (bursty) and 38ms (Poisson) at PLR $< 10^{-6}$. They also investigated wavelength re-use in the wavelength routing network as a function of the lightpath setup overhead and showed that for good wavelength reuse, the overhead time must be less than about 10ms, which may not be possible for networks which span large distances (e.g. national or transcontinental).

Finally, a novel idea is proposed by Detti et al [40] called Optical Composite Burst Switching (OCBS) in which if a burst cannot be carried on an outgoing wavelength at a particular switch, it is broken up into two parts, one which is blocked and one which is not. The rationale is that bursts are asynchronous and vary in length, and their start and end times are not aligned, so the arriving burst may not be fully blocked by the incumbent burst. The part which is not blocked can continue to the destination and still be useful, since it should contain a number of useable IP packets still. This relies on the switches being able to distinguish the IP packets within the burst. Their method uses a framing format and synchronisation on a header CRC, similar to ATM and HDLC, however this method is bitrate and modulation format dependent and so not suitable for transparent optical networks. If, however, a method could be found to delineate IP packets photonically, this method could certainly be used. Detti et al present an analysis based on Horizon scheduling with a Markov chain formulation to calculate loss probabilities of packets in the burst, which has not been seen to date and could be extended to traditional JET and JIT networks.

4. Work Completed and In Progress

This research has identified burst loss ratio as a key problem to be overcome in OBS. I have developed a simulation framework in the Opnet™ Modeler simulation package, for simulating OBS networks with arbitrary topologies and system parameters. I have developed simulations for standard OBS and the offset time-based QoS of Qiao and Yoo [24]. The results

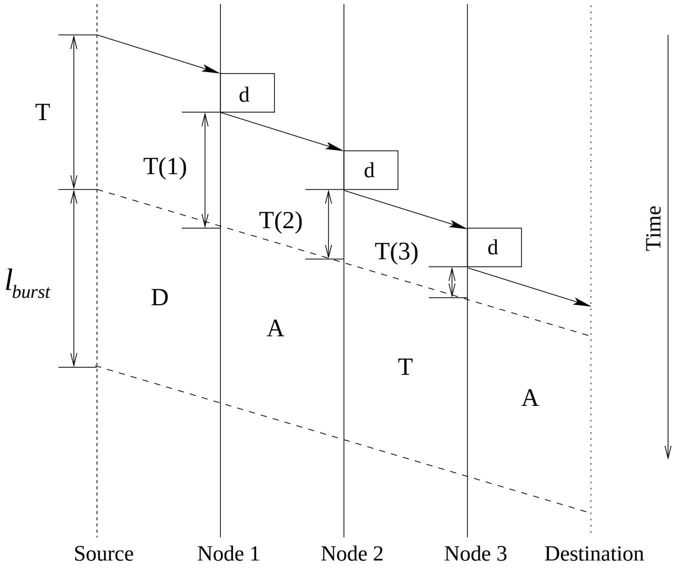

FIGURE 1. Burst and control packet propagation in OBS. The gradual reduction of residual offset times is shown.

of these simulations have proved to be well matched to those of Qiao and Yoo (see SECTION 4.2).

4.1. **Merit-based Quality-of-Service.**

As each burst traverses the network, its residual offset time $T(i)$ decreases at each subsequent hop i. $T(i)$ represents the original offset time T reduced by a processing time at each hop: $T(i) = T - (i-1)d$ where d is the control packet processing time at each switch, as shown in Fig. 1. Therefore a burst which has travelled further is more likely to be dropped than a burst which has only just begun its journey. The question arises: can we give the longer travelled burst a helping hand and at the expense of the shorter travelled burst? The answer is yes since the JET/OBS offset time T allows us to see into (and alter) "the future", or more correctly the set of bursts which have already been accepted by the switch for future transmission. If we decide that a new burst B_1 is more deserving of the transmission bandwidth than a burst B_2 which has been accepted already, then we can eject B_2 from the schedule at any point up until its arrival at the switch. We must then signal the downstream nodes to eject B_2 from their schedules as well, to recover bandwidth which could still be used.

The question remains as to how to decide which bursts are more deserving of bandwidth. We introduce a measure of burst merit $V(\mathbf{x}_n)$ where $\mathbf{x}_n$ is a vector representing the characteristics of the burst B_n. This vector may comprise information such as:

- Residual offset time, R_n.
- Original offset time, T_n.
- Number of hops traversed, H_n.
- Length of the intended source-to-destination route, L_n.
- Burst length, B_n.

– Class of service, C_n.

Here R_n is the value of $T(i_n)$ for the nth burst, assuming that the nth burst is at its i_nth hop. Thus $\mathbf{x}_n$ could be written as:

$$\mathbf{x}_n = \left[\begin{array}{cccccc} R_n & T_n & H_n & L_n & B_n & C_n \end{array} \right]. \tag{1}$$

An appropriate definition for $V(\mathbf{x}_n)$ would seek to give preference to bursts which have already used a significant transmission bandwidth (i.e. those which have traversed more hops). It would also seek to give bursts with further distance to travel more preference, since they are more likely to be blocked (there are more switches at which they can be blocked), wasting their transmission bandwidth. Given that a low $T(i)$ causes higher blocking, we might also seek to give preference to bursts with lower $T(i)$. Finally, higher priority service classes should not be preempted by lower priority classes, in a network which offers multiple grades of service.

The choice of $V(\mathbf{x}_n)$ is in some senses arbitrary, and it is a goal of this research to explore different $V(\mathbf{x}_n)$ mappings and their impact on network burst loss ratio performance. As a first cut at the problem, the following was proposed:

$$V(\mathbf{x}_n) = C_n + (1 - \exp(-\varphi(\mathbf{x}_n))), \tag{2}$$

where $\varphi(\mathbf{x}_n)$ is to be determined and C_n is the class of service component of $\mathbf{x}_n$. Suppose burst B_1 is represented by $\mathbf{x}_1$ and burst B_2 by $\mathbf{x}_2$, and that $\mathbf{x}_1$ has class of service component C_1 and $\mathbf{x}_2$ has class of service component C_2. The function $V(\mathbf{x}_n)$ has the property that if $\mathbf{x}_1$ and $\mathbf{x}_2$ are such that $C_1 < C_2$, then $V(\mathbf{x}_1) < V(\mathbf{x}_2)$, so that the high priority burst is never preempted by a lower priority burst. This is due to the asymptotic behaviour of the $1 - \exp(-\varphi(\mathbf{x}_n))$ term. Since this term is monotonically increasing, $\varphi(\mathbf{x}_n)$ determines the ordering of bursts within each class. As an initial attempt, the following $\varphi(\mathbf{x}_n)$ was chosen:

$$\varphi(\mathbf{x}_n) = H_n \cdot L_n. \tag{3}$$

This is simple and should give preference to bursts which have already come a long way (large H_n) and have a long intended path (large L_n), as desired. $V(\mathbf{x}_n)$ is used to induce an order on all bursts. We introduce the relation $<$ on bursts, such that if bursts B_1 and B_2 are described by vectors $\mathbf{x}_1$ and $\mathbf{x}_2$ then

$$B_1 < B_2 \quad \text{iff} \quad V(\mathbf{x}_1) < V(\mathbf{x}_2). \tag{4}$$

I call this “merit-based” quality-of-service (QoS), since it rates the merit of each burst in deciding which should be given assistance in reaching their destinations.

4.2. **Simulation of Merit-based Quality-of-Service.**

I have developed a simulation which implements this idea. When a new burst B arrives at a switch, the LAUC-VF algorithm is executed to attempt to find an outgoing wavelength for it. If one cannot be found, the wavelength set is scanned and the bursts ranked along side B. The burst with least merit, B_l, is ejected from the schedule to allow B to be inserted. Then the downstream switches are signalled to cancel their reservations for B_l. If there is more than one burst that would need to be ejected on a given wavelength to accomodate B, then currently the one with the highest merit is used to represent the group. It might perhaps be better to use a sum of the merit $V(\mathbf{x}_j)$ for every burst in the group, weighted by the size of each burst. This will be studied in future along with other variations.

Results from the simulation are shown below for a 5×5 mesh-torus network topology. All results were for a network with 8 wavelengths each of 10G/s capacity, with a ninth control wavelength at 2.5G/s. All nodes have equal distance giving a 500μs propagation delay. The control packet processing delay at each switch is $0.1ms$ and the switch is assumed to require 0.1μs to set up a connection from an input port to an output port. There are two classes in the system, each contributing the same amount of load, and they both have the same burst length distribution (exponential) with the same average burst length of 50,000 bytes. Arrivals are according to a Poisson process. For the offset time-based QoS scheme, the offset time difference between classes is 120μs, or 3 times the average burst length. Each point on the following graphs is the average of 30 samples of the same network with different random number generator seeds for each sample. Each sample run of the network lasted 20ms of simulated time, which was seen to give steady state behaviour to all statistics.

Figure 2 shows the burst loss ratio for the merit-based QoS scheme. The burst loss ratio is significantly lower for the class 1 traffic than for in the equivalent network for with the Qiao and Yoo method (offset time-based). The results for the offset time-based scheme for the same network are shown in Fig. 3. In fact, even the combined loss ratio is reduced, whereas in Qiao

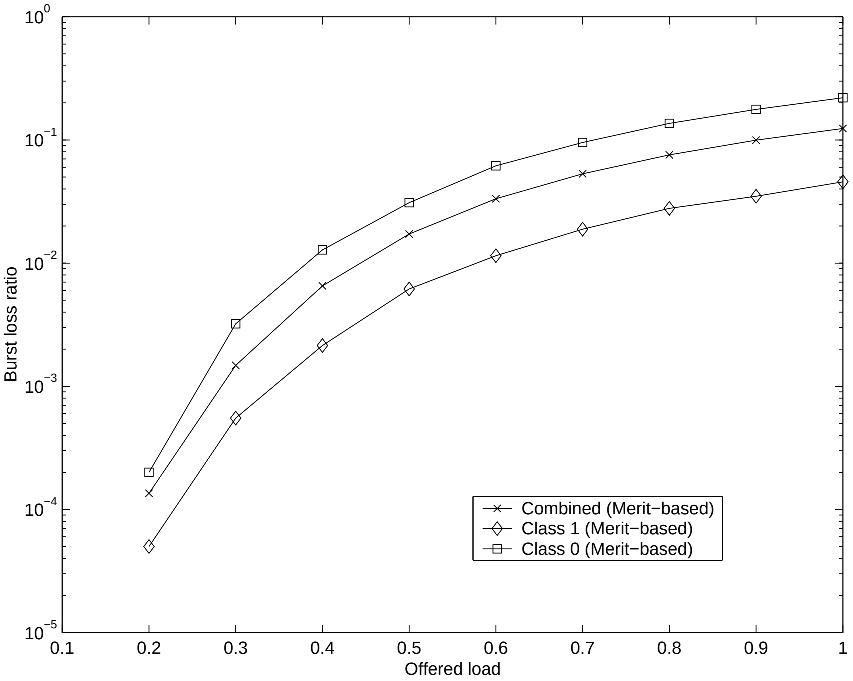

FIGURE 2. Burst loss ratio vs offered load for class 0, class 1 and combined traffic in Merit-based QoS.

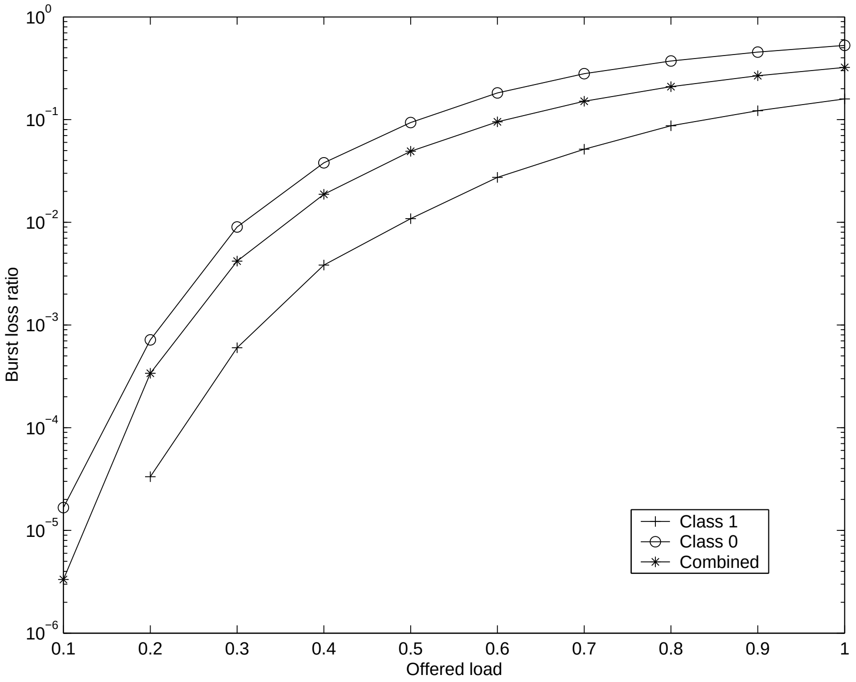

Figure 3. Burst loss ratio vs offered load for class 0, class 1 and combined traffic in Offset time-based QoS.

and Yoo's method it remains unchanged. Also, the class 0 traffic in the Merit-based QoS scheme actually fares better than the combined loss ratio of the Offset time-based scheme. As a point of comparison, Fig. 4 shows the class 1 loss ratio for both Merit-based QoS and Offset time-based QoS on the same graph. This graph shows that Merit-based QoS outperforms Offset time-based QoS by a significant margin for high offered load, however at low offered load (less than 0.3) the Offset time-based QoS performs better. It is not clear why this is so. Also, the gains of merit-based QoS steadily increase with offered load, such that at load of 1.0 the loss ratio is roughly a quarter of the Offset time-based scheme (0.0457 vs 0.1592). The decrease in overall loss ratio for merit-based QoS is further reflected in the utilisation for the network (Fig. 5). The merit-based QoS mechanism yields significantly higher utilisation for high offered loads, and almost the same performance when the load is low, since at low load the burst ejection mechanism would be rarely invoked because LAUC-VF should be able to find a free wavelength most of the time.

4.3. OBS as a Queueing System Optimisation Problem.

In addition to the simulation work described above, I have also collaborated with Hai Le Vu and Moshe Zukerman (see Section 6) on an optimisation problem which investigates the capacity of burst switching systems with regard to the speed of the control processor of the burst switch. The question posed is: what is the maximum throughput attainable in an OBS system if we require control packet queueing to be negligible and wish to limit the burst loss ratio? We ask this question because it is desirable to choose an

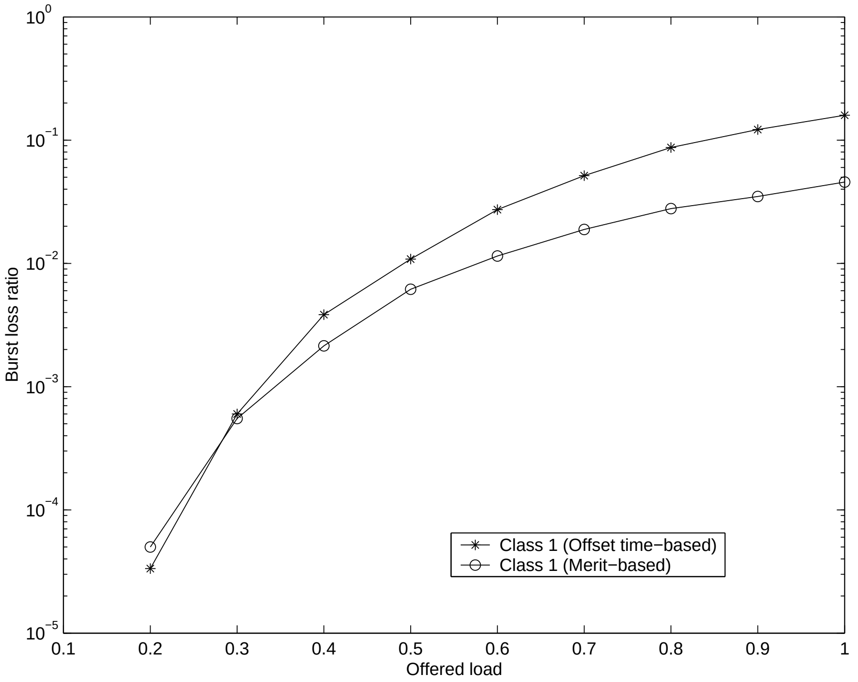

FIGURE 4. Burst loss ratio vs offered load for class 1 traffic in both Merit-based and Offset time-based QoS.

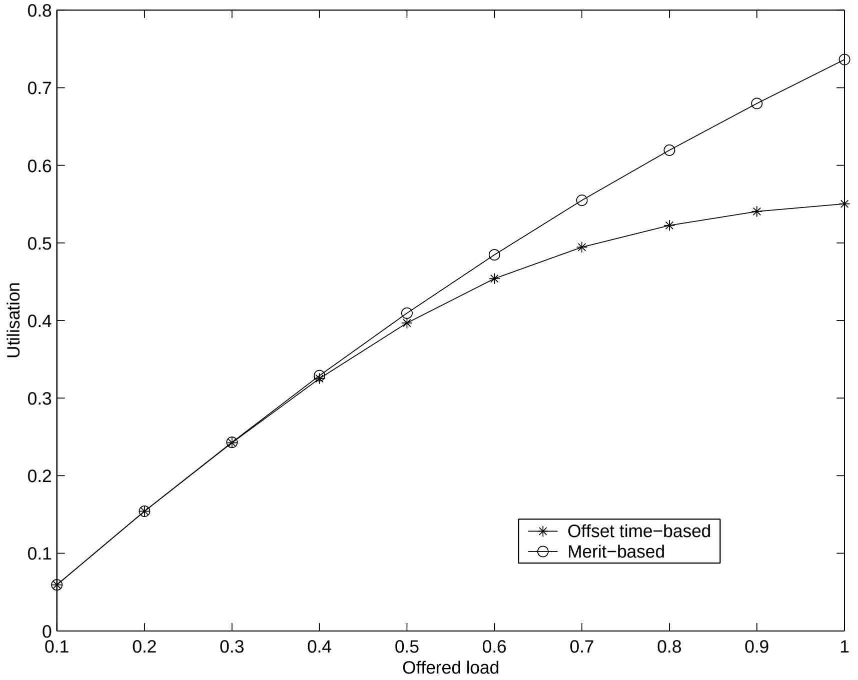

FIGURE 5. Network utilisation vs offered load.

average burst length B which does not result in too many control packets being sent. If control packets congest a control processor, they must be queued and therefore their waiting time becomes non-deterministic. OBS requires strict timing of the burst relative to the control packet, so this is not desirable. We model an output port of a switch as an $M/G/k/k$ queuing system, and the control processor as a $G/G/1$ queue, and we place the following constraints on the system:

$$\Pr\{Q > 0\} = \frac{\lambda}{\mu} \leq \epsilon_1 \tag{5}$$

$$E_k(A(B)) \leq \epsilon_2. \tag{6}$$

Here, λ is the arrival rate of control packets and bursts to the output port, μ is the service rate of control packets, $A(B) = \lambda B$ is the throughput of the data channels and $E_k(A)$ is the Erlang loss probability for an $M/G/k/k$ system with offered traffic A. Q is the length of the control processor queue. The parameters $\epsilon_1, \epsilon_2 > 0$ are chosen by the designer to ensure minimal queueing in the control processor (5) and minimal loss of bursts in the data path (6).

It was found that throughput $A(B)$ is maximised at

$$A_{\max} = E_k^{-1}(\epsilon_2), \tag{7}$$

which is achieved for

$$B \geq \frac{1}{\mu\epsilon_1} E_k^{-1}(\epsilon_2). \tag{8}$$

($E_k^{-1}(\epsilon_2)$ is the throughput required for an $M/G/k/k$ queueing system to have a blocking probability of ϵ_2).

Traffic modelling showed that an appropriate choice of B is dependent on traffic conditions and the burst assembly process, specifically the number of sources N, the maximum allowable burst assembly delay t^*, and the maximum allowable burst size $B_{\max}$, together with the traffic intensity per source, $A(B)/N$.

The system is feasible if $(N/t^*) \leq \mu\epsilon_1$. This work has been submitted as a paper to *IEEE Communications Letters*.

5. Research Plan and Time-line

5.1. **Immediate Plans.**

Immediate plans for continuation of this research will focus on finding alternative measures of burst merit $V(\mathbf{x}_n)$ and investigating their impact on network performance. The overarching theme of this research is overcoming the effects of the shrinking residual offset time. With $V(\mathbf{x}_n)$ defined as in (2), possible definitions of $\varphi(\mathbf{x}_n)$ for further investigation are:

$$\begin{aligned}
\varphi(\mathbf{x}_n) &= \frac{H_n L_n}{(L_n - H_n)} \quad \text{(boost for bursts close to destination)}\\
\varphi(\mathbf{x}_n) &= \frac{1}{R_n} \quad \text{(boost for small offset)}\\
\varphi(\mathbf{x}_n) &= \frac{T_n}{R_n} \quad \text{(boost for small relative offset).}
\end{aligned}$$

All of these could also be weighted by the burst length B, and linear combinations of these should also be considered. Simulations will also be run to record the distribution of residual offset times and the burst loss ratio as a function of residual offset time, to see how the above algorithm affects them. The results presented in SECTION 4 were for a 5×5 mesh-torus network with two classes of service and 8 wavelengths per fibre. Qiao and Yoo showed that as more classes are added to the network, in their QoS scheme the loss ratio of the highest priority class is improved dramatically. In some sense it profits from the added contention amongst the lower priority classes. To investigate whether similar behaviour affects the merit-based QoS scheme, the simulations will be re-run with more classes of service. Similar comparative simulations will be run with more data wavelengths per channel to examine how increased transmission resources affect the burst loss ratio.

Because the mesh-torus topology is highly connected and has relatively short routes for the number of nodes in the network, it probably gives a more optimistic view of network performance. To this end, further simulations will be run using both the Qiao and Yoo QoS scheme as well as the merit-based QoS scheme for a wider range of topologies. These will include the ring topology, the partially meshed ring topology with varying degrees of connectivity and a linear bus topology (where the OBS switch functions as an Add-Drop Multiplexer). The NSFNET topology is a standard topology used in the literature to model a North American network which will also be studied in this research. It would be interesting also to try to evaluate a sparsely connected, largely acyclic network which could model a national network for Australia's main population centres also.

Along with less connected and less regular topologies it would be instructive to consider asymmetric traffic loads, since most Internet traffic tends to be asymmetric both between the two sides of a bidirectional session (e.g. HTTP traffic) and also in terms of distribution over the network – in the Internet there are a relatively small number of nodes providing different services to a much larger number of clients. This has the potential to create significant bottlenecks and congestion on critical links in the transport network with which OBS, being essentially bufferless, may be ill-equipped to cope.

A related issue is modelling of traffic statistics. Poisson arrival statistics have been used in the simulations presented in SECTION 4, however this will be extended to self-similar, heavy tailed models such as Pareto, M/Pareto and Fractional Gaussian Noise, as in [31, 38].

5.2. **Alternate Switch Architectures.**

A major failing of OBS is the problem of the reduction in residual offset time $T(i)$ as the burst traverses the network, which induces large loss ratios. Whilst most other techniques, including the merit-based QoS method presented above, attempt to make the most of a bad situation to reduce blocking probability, this ignores the potential which might exist to explore new switch architectures which could enable the offset time to be obviated, or equalised across all arrivals. A major area of investigation for this research will be to focus on this issue and attempt to find novel uses for photonic technologies to assist in eliminating the residual offset time as a difficulty.

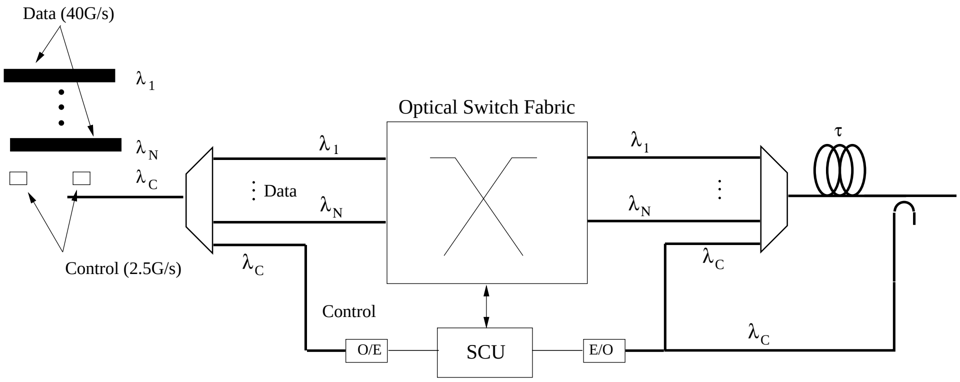

FIGURE 6. OBS switch with output delay providing capability to selectively increase the offset time by amount τ seconds, or decrease it to 0.

One proposal is to place a fibre delay line on the output of the switch, equivalent to a τ second delay, with a short fibre used to couple the output of the switch control unit (SCU) to the fibre *after* the delay line, see Fig. 6. In the figure only one input and output port are shown, for clarity. In this configuration, the SCU, which is assumed to convert the incoming control packets into electrical form, can buffer the packet and time its release to give a precise offset time between the burst and the control packet. If it could also release the packet into the fibre before the FDL then it would be possible to provide both decreases and *increases* in the offset time. Investigating this proposal and how best to use this "manageable" offset time will be a significant problem to be solved. An alternative to this approach is to deviate from traditional offset-time oriented, decoupled-control OBS by having bursts preceded by their control packets with no offset time and on the same wavelength. An offset time can be provided by placing an FDL between the inlet and the switching matrix, but by tapping off a small amount of light before the FDL (using e.g. a 10/90 splitter), the control information for each burst (probably modulated at a lower bit rate) can be read electronically and fed directly into the SCU with no delay (Fig. 7 – again only one port is shown for clarity). Supposing the delay line is equivalent to a τ second delay again, all bursts now have a τ second delay between their control information being processed and the data payload entering the switch. The control information is re-written onto the output fibre to be again immediately in front of the burst data on the same wavelength. Normal optical packet switching makes τ just long enough for the header to be processed, however by providing for instance $\tau = 100\mu$s we have effectively given a significant degree of extra freedom to the wavelength scheduler. Novel scheduling algorithms and control packet queueing disciplines will be investigated under this scheme to enable evaluation of its ability to improve the blocking probability of OBS.

These schemes are also key to fulfilling the stated objective of this research to explore the relationships that exist between OBS and optical packet switching.

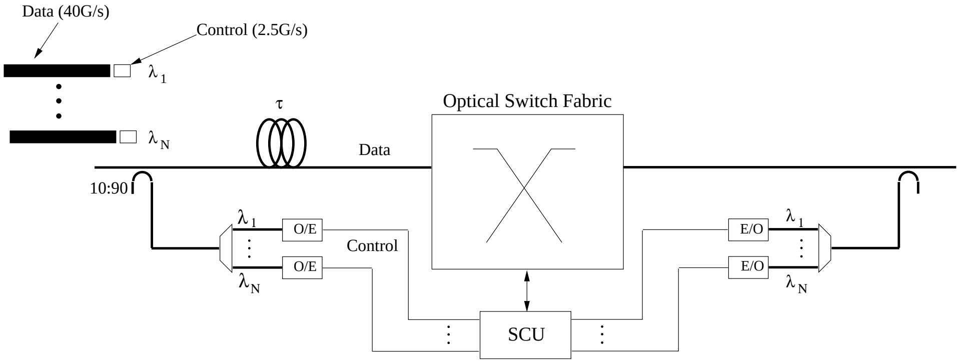

Figure 7. OBS switch with input delay providing offset time of τ seconds. The data payload and control header are sent on the same wavelength channel, and immediately follow one another. They are recombined at the output coupler in the same manner.

5.3. Technology aspects.

Optical technologies behave significantly differently from electronic technologies. To date researchers have concentrated on the logical properties of OBS, without considering the impact the underlying optical and photonic technologies may have. Two areas have been identified as suitable for further investigation in this regard. There has been no published work to date on the impact of limited range wavelength conversion in OBS. All OBS architectures assume wavelength conversion capability from any input wavelength to any output wavelength which in reality is not the case. Yates et al conducted a study of limited range conversion in [14] for wavelength routed networks. It is proposed that similar analysis be done under the OBS regime. Both extension of Yates' analysis where possible and development of simulation models will be considered.

There has been some work published on the effect of Erbium Doped Fibre Amplifier (EDFA) gain saturation on power levels in fibre carrying bursty traffic [41]. Since OBS is intended to carry bursty traffic, it would be interesting to perform similar analysis and simulation of how EDFAs react to bursty traffic that has been shaped by losses in an OBS network. Using the already developed simulations to capture traffic traces which could be fed into a transmission system simulation tool such as VPITransmissionMaker is an avenue for seeing how much of a problem this could be.

5.4. Flow controlled OBS.

A key aspect of OBS which has not been investigated to date is its affects on the higher level protocols which it will carry. For example, the most likely network protocol for OBS to carry is the Internet Protocol (IP). The IP protocol carries Transmission Control Protocol (TCP) as one possible and popular end-to-end transport protocol, however TCP is a flow-controlled protocol which uses packet loss as a signal of network congestion to adjust its rate. However, this assumes that the loss is in the IP routers along the path, with the links between routers assumed to be nearly ideal (no loss). In a wavelength routed or circuit switched underlying transport network this would be the case, however OBS introduces the additional loss in the transport network. The question arises: how will this additional loss affect the behaviour (loss, fairness, stability) of TCP, and would it be necessary for routers or sources to adjust their behaviour to compensate? If changes to routers could compensate then this would be preferable to having to change TCP itself, which is installed in a vast number of computers compared to the number of routers.

Taking this one step further, is it possible for OBS switches themselves to participate in the TCP flow-control process to mitigate undesired behaviour? Might it be better to execute a second flow control protocol within the OBS network itself rather than interacting with TCP? Maybe the active queue management (AQM) algorithm used by burst assemblers should be different from typical IP routers to account for the idiosyncracies of OBS.

Simulations will be constructed to try to answer these questions, together with analysis based on traditional flow-control analysis, such as max-min fairness [15] and the macroeconomic model [42].

5.5. **Analytic Modelling.**

It is expected that the work in the above sections will be carried out through a combination of simulations, principally using the OPNET™program, together with appropriate analytical work. Yates et al [14] may provide insight and a basis for considering the effect of limited wavelength conversion in OBS, whilst Dolzer et al [30] and Detti et al [40] give analyses based on queueing theory and Markov chains which could be extended to help analyse the merit-based QoS scheme and other approaches outlined as future work presented above.

5.6. **Timeline.**

A timeline for the remainder of my PhD candidature is shown in the table overleaf. The time to be allotted for the work outlined above is shown in this table.

Date	Work Description
3/2002 – 4/2002	Extend Merit-based QoS simulation to investigate different measures of burst merit and different network topologies. Explore the dependency of loss ratio on number of wavelengths per fibre, number of service classes, size of network, non-uniform traffic matrix and self-similar traffic models. Use simulation to quantify arrival offset time distribution and blocking probability of arrivals vs offset time. Develop analytical models to describe loss behaviour of Merit-based QoS.
5/2002 – 7/2002	Investigate output-buffered optical burst switch architecture. Develop control packet output scheduling algorithms to improve loss ratio performance over standard OBS.
8/2002 – 11/2002	Investigate input-buffered optical burst switch architecture. Develop simulation models to implement novel control packet queueing disciplines to control how and when each burst is given a place in the wavelength assignment schedule on an outgoing port. Attempt to use such algorithms in a similar manner to Merit-based QoS outlined above, to selectively give priority to bursts which are more deserving of transmission bandwidth
12/2002 – 2/2003	Attempt to develop analytical models of performance of algorithms developed for the input and output buffered switch architectures. This will depend on the performance of the architectures shown in simulation. It may be worthwhile concentrating only on the one which performed better.
3/2003 – 4/2003	Attempt to incorporate limited degree wavelength converter capabilites and limited converter availability into simulation models for input- and output-buffered switch architectures and assess the impact on the overall viability of OBS. Is OBS viable if this key technology is not as flexible as the ideal case? Attempt to provide analysis to explain the behaviour seen in simulation
5/2003 – 9/2003	Develop simulations to investigate the behaviour of Transmission Control Protocol (TCP) when carried over an underlying OBS network, in particular the impact of TCP end-to-end flow control. Examine TCP stability, loss, throughput and congestion. Attempt to explain observed behaviour using analytic modelling. Extend models (simulation and analysis) to incorporate active queue management (AQM) of control packet queues and explicit congestion notification.

Date	Work Description (cont'd)
10/2003 – 11/2003	Capture traces of line traffic from OBS simulations for different networks and algorithms studied above. Use these traces as input to a transmission network simulation tool (e.g. VPITransmissionMaker) to investigate the effects of bursty traffic on power transients in EDFAs and other physical layer optical and photonic components.
12/2003 – 2/2004	Complete thesis.

6. Published Work

- Jolyon White, Moshe Zukerman and Hai Le Vu. "A Framework for Optical Burst Switching Network Design," submitted to *IEEE Communications Letters*.

References

[1] L. G. Roberts, “Beyond Moore’s Law: Internet growth trends,” *Computer*, vol. 33, pp. 117–119, Jan. 2000.

[2] V. Paxson and S. Floyd, “Wide area traffic: The failure of Poisson modelling,” *IEEE Trans. Networking*, vol. 3, pp. 226–244, Jun 1995.

[3] A. E. Willner, “Mining the optical bandwidth for a terabit per second,” *IEEE Spectrum*, vol. 34, pp. 32–41, Apr. 1997.

[4] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*. San Francisco, CA, USA: Morgan Kaufmann Publishers, Inc, 1998.

[5] B. Mukherjee, “WDM-based local lightwave networks part I: Single-hop systems,” *IEEE Network*, vol. 6, pp. 12–27, May 1992.

[6] B. Mukherjee, “WDM optical communication networks: Progress and challenges,” *IEEE J. Selected Areas in Communications*, vol. 18, pp. 1811–1824, Oct. 2000.

[7] S. Okamoto, “Photonic transport network architecture and OA&M technologies to create large-scale robust networks,” *IEEE J. Selected Areas in Communications*, vol. 16, pp. 995–1007, Sept. 1998.

[8] E. Karasan and E. Ayanoglu, “Performance of WDM transport networks,” *IEEE J. Selected Areas in Communications*, vol. 16, pp. 1081–1096, Sept. 1998.

[9] I. Chlamtac, A. Ganz, and G. Karmi, “Lightpath communications: An approach to high bandwidth optical WAN’s,” *IEEE Trans. Commun.*, vol. 40, pp. 1171–1182, July 1992.

[10] A. Birman, “Computing approximate blocking probabilities for a class of all-optical networks,” *IEEE J. Selected Areas in Communications*, vol. 14, pp. 852–857, June 1996.

[11] R. A. Barry and P. A. Humblet, “Models of blocking probability in all-optical networks with and without wavelength changers,” *IEEE J. Selected Areas in Communications*, vol. 14, pp. 858–867, June 1996.

[12] M. Kovačević and A. Acampora, “Benefits of wavelength translation in all-optical clear-channel networks,” *IEEE J. Selected Areas in Communications*, vol. 14, pp. 868–880, June 1996.

[13] S. Xu, L. Li, and S. Wang, “Dynamic routing and assignment of wavelength algorithms in multifiber wavelength division multiplexing networks,” *IEEE J. Selected Areas in Communications*, vol. 18, pp. 2130–2137, Oct. 2000.

[14] J. M. Yates, J. P. R. Lacey, M. P. Rumsewicz, and M. A. Summerfield, “Performance of networks using wavelength converters based on four-wave mixing in semiconductor optical amplifiers,” *IEEE J. Lightwave Tech.*, vol. 17, pp. 782–791, May 1999.

[15] D. Bertsekas and R. Gallager, *Data Networks*. Upper Saddle River, New Jersey: Prentice Hall, 2nd ed., 1992.

[16] S. Yao, S. J. B. Yoo, B. Mukherjee, and S. Dixit, “All-optical packet switching for metropolitan area networks: Opportunities and challenges,” *IEEE Commun. Mag.*, vol. 39, pp. 142–148, Mar. 2001.

[17] R. S. Tucker and W. D. Zhong, “Photonic packet switching: An overview,” *IEICE Trans. Electron.*, vol. E82-C, pp. 202–212, Feb. 1999.

[18] J. Spring, R. M. Fortenberry, and R. S. Tucker, “Photonic header replacement for packet switching,” *Electron. Lett.*, vol. 29, pp. 1523–1524, Aug 1993.

[19] I. Chlamtac, A. Fumagalli, and C.-J. Suh, “Multibuffer delay line architectures for efficient contention resolution in optical switching nodes,” *IEEE Trans. Commun.*, vol. 49, pp. 2089–2098, Dec. 2000.

[20] I. Chlamtac, A. Fumagalli, and C.-J. Suh, “Switching multi-buffer delay lines for contention resolution in all-optical deflection networks,” in *Proceedings of GLOBECOM 96*, vol. 3, pp. 1624–1628, 1996.

[21] M. Yoo and C. Qiao, “Just-Enough-Time (JET): A high speed protocol for bursty traffic in optical networks,” in *IEEE/LEOS Technologies for a Global Information Infrastructure*, pp. 26–27, Aug. 1997.

[22] M. Yoo and C. Qiao, "A new optical burst switching protocol for supporting quality of service," in *Proc. SPIE'98 Conf. All Optical Commun. Syst.: Architecture, Control, Network Issues*, vol. 3531, (Boston), pp. 396–405, Nov. 1998.

[23] G. C. Hudek and D. J. Muder, "Signalling analysis for a multi-switch all-optical network," in *Proceedings of IEEE International Conference on Communications, ICC 95*, pp. 1206–1210, 1995.

[24] M. Yoo, C. Qiao, and S. Dixit, "QoS performance of optical burst switching in IP-Over-WDM networks," *IEEE J. Selected Areas in Communications*, vol. 18, pp. 2062–2071, Oct. 2000.

[25] C. Qiao and M. Yoo, "Choices, features and issues in optical burst switching," *Optical Networks Magazine*, vol. 1, pp. 36–44, Apr. 2000.

[26] C. Qiao and M. Yoo, "Optical burst switching (OBS) - a new paradigm for an optical internet," *J. High Speed Networks*, vol. 8, pp. 69–84, Jan. 1999.

[27] J. Turner, "Terabit burst switching," *J. High Speed Networks*, vol. 8, pp. 3–16, Jan. 1999.

[28] J. Wei, J. Pastor, R. Ramamurthy, and Y. Tsai, "Just-in-time optical burst switching for multiwavelength networks," in *IFIP Broadband Commun.*, (Hong Kong), pp. 339–352, Nov. 1999.

[29] J. Wei and R. McFarland, "Just-in-time signaling for WDM optical burst switching networks," *IEEE J. Lightwave Tech.*, vol. 18, pp. 2019–2037, Dec. 2000.

[30] K. Dolzer, C. Gauger, J. Späth, and S. Bodamer, "Evaluation of reservation mechanisms for optical burst switching," *Int. J. Electron. Commun.*, vol. 55, no. 1, pp. 18–26, 2001.

[31] Y. Xiong, M. Vandenhoute, and H. Cankaya, "Control architecture in optical burst switched WDM networks," *IEEE J. Selected Areas in Communications*, vol. 18, pp. 1838–1851, Oct. 2000.

[32] M. Yang, S. Q. Zheng, and D. Verchere, "A QoS supporting scheduling algorithm for optical burst switching DWDM networks," in *Proceedings of GLOBECOM 01*, pp. 86–91, 2001.

[33] X. Wang, A. Saito, H. Morkawa, and T. Aoyama, "Distributed wavelength assignment algorithm for optical bursts in WDM networks," in *Proceedings of OECC/IOOC 2001 Conference on Optical Communications*, (Sydney), pp. 191–193, 2001.

[34] L. Xu, H. Perros, and G. Rouskas, "Techniques for Optical Packet Switching and Optical Burst Switching," *IEEE Commun. Mag.*, pp. 136–142, Jan. 2001.

[35] C. Qiao, "Labeled optical burst switching for IP-over-WDM integration," *IEEE Commun. Mag.*, vol. 38, pp. 104–114, Sept. 2000.

[36] F. Callegati, H. C. Cankaya, Y. Xiong, and M. Vandenhoute, "Design issues of optical ip routers for internet backbone applications," *IEEE Commun. Mag.*, vol. 37, pp. 124–128, Dec. 1999.

[37] M. Yoo, C. Qiao, and S. Dixit, "Optical burst switching for service differentiation in the next-generation optical internet," *IEEE Commun. Mag.*, vol. 39, pp. 98–104, Feb. 2001.

[38] A. Ge, F. Callegati, and L. Tamil, "On optical burst switching and self similar traffic," *IEEE Commun. Lett.*, vol. 4, pp. 98–100, Mar. 2000.

[39] M. Düser and P. Bayvel, "Performance of a dynamically wavelength-routed optical burst switched network," *IEEE. Photonic Technology Lett.*, vol. 14, pp. 239–241, Feb. 2002.

[40] A. Detti, V. Eramo, and M. Listanti, "Performance evaluation of a new technique for ip support in a WDM optical network: Optical composite burst switching (OCBS)," *IEEE J. Lightwave Tech.*, vol. 20, pp. 154–165, Feb 2002.

[41] M. Karásek, L. A. Rusch, and M. Menif, "Suppresion of output power and NF excursions in cascades of highly inverted EDFAs with packet switched traffic," *Fiber and Integrated Optics*, vol. 20, pp. 269–277, 2001.

[42] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: shadow prices, proportinal fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, Mar. 1998.