WDM Bandwidth Allocation

AAPN Technical Report 2004-03

Nahid Saberi, Mark Coates
Department of Electrical and Computer Engineering
McGill University

June, 2004

1 Introduction

In an all-optical network, wavelength division multiplexing (WDM) has been introduced as a means to support the rapidly increasing bandwidth demand of the network users. Partitioning the enormous bandwidth in optical networks into WDM channels makes the optical bandwidth compatible with the speed of electronic components. In WDM systems the end users communicate via optical channels called lightpaths. In networks with *continuity constraint* a lightpath is required to occupy the same wavelength on all fiber links, while in *wavelength convertible* networks a lightpath can occupy different wavelengths on different fiber links along the path. Wavelength conversion allows networks to support more lightpaths, but it is expensive and difficult to implement [1]. Therefore in this study we mainly consider networks with continuity constraint.

In these networks, resource allocation is a key problem which is addressed in many different contexts, including routing and wavelength assignment (RWA), optical burst/packet switching, photonic slot routing, TDM/WDM assignment, and broadcast-and-select technology.

Optical backbone networks multiplex a large amount of traffic coming from numerous users on circuit-switched wavelength paths. This technology (the so-called wavelength routing approach) has been widely studied in the literature [1]. In contrast, wavelength routing in access and metro networks with a reduced level of traffic aggregation is not an adequate solution [2]. In these areas, other approaches such as optical burst/packet switching and time-division multiplexing over WDM channels provide a more dynamic bandwidth allocation. Based on TDM/WDM technology several different approaches have been developed such as photonic slot routing and TDM/WDM bandwidth reservation in broadcast and select networks.

In this article we survey the background literature and recent developments in WDM resource allocation techniques. This information provides the basis for our research which deals with resource allocation and bandwidth sharing in a particular type of network with a star topology referred to as AAPN[1] [3]. The next section introduces routing and wavelength assignment in WDM networks. Section 3 presents an overview of recent optical switching deployments. Section 4 provides an extensive description of bandwidth reservation in WDM broadcast-and-select networks.

2 Routing and Wavelength Assignment Schemes

The problem of assigning wavelengths and paths to a set of requests for bandwidth between source-destination pairs in WDM networks is referred to as the routing and wavelength assignment (RWA) problem. It has been proved that the RWA problem is NP[2]-Complete[3] [5], and partitioning this problem into two subproblems, (i) routing, and (ii) wavelength assignment, makes it more tractable. Numerous solutions for either of the problems have been proposed; these can be classified under static and dynamic approaches. Depending on the traffic pattern the problem can be formulated differently. When the traffic pattern is static, the requests are fixed for a long duration of time as the connections are long-lived. Therefore, the entire demand is known in advance (off-line information). Once the connections are established there is no need for further operations. In a dynamic traffic pattern the requests arrive on-line and may depart after a while. Therefore the RWA problem has to

[1] Agile All-Photonic Network

[2] Nondeterministic polynomial-time problem; a problem is said to be NP if there exists a nondeterministic polynomial-time algorithm that recognizes the elements of this problem (i.e. that we can test in polynomial-time whether a possible solution is a solution).

[3] A problem in NP is said to be NP-Complete if finding a deterministic polynomial-time algorithm for solving this problem allows us to solve any problem in NP in polynomial-time (for more information regarding the theory of NP-completeness refer to [4]).

consider the dynamics of the related parameters for assigning paths and wavelengths to such traffic. For the routing subproblem there are three basic approaches known as fixed routing, fixed alternate routing and adaptive routing. For the wavelength assignment component many approaches have been proposed so far such as First-Fit, Least-Used, Most-Used, Min-product, Least loaded, Max-SUM and Wavelength Reservation. We now present a review of various routing and wavelength assignment approaches for both off-line and on-line cases; the review is partially derived from [1].

2.1 Routing Schemes

Many routing algorithms has been proposed in literature which can be classified in three basic groups:

2.1.1 Fixed Routing

Fixed routing, the simplest way of performing routing in a network, chooses a fixed path from a node to each destination. Fixed shortest-path routing is an instance of the fixed routing approach, which can be implemented by using standard shortest path algorithms, such as Dijkstra's algorithm or the Bellman-Ford algorithm. The major problem with this approach arises when the resources are insufficient to meet the demand, which tends to a high blocking probability [1].

2.1.2 Fixed-Alternate Routing

Fixed alternate routing constructs a list of options for routing a path between each source-destination pair. The algorithm may use different criterions for constructing and sorting the list, such as a shortest-path criterion or the load at each link. For example, the list may include the first, the second and the third shortest paths between each source-destination pair. This routing approach reduces the blocking probability compared to fixed routing. In addition, in the case of path failure, it can provide some degree of flexibility for rerouting the connections [1].

2.1.3 Adaptive Routing

Adaptive routing is a routing algorithm that determines the available paths across a network and based on the network state evaluates them and chooses the one that will provide the best path for a connection. This approach works well for networks with dynamic states. Examples of adaptive routing include:

Adaptive Shortest-Cost-Path. This approach assigns a cost to each link in the network. The cost is determined based on the state of the link in the network. Each unused link is assigned a cost of 1 unit, each used link is assigned a cost of ∞, and each used link with a wavelength convertor is assigned a cost of c units, where the value of c is defined in such a way as to avoid the use of wavelength convertors as long as any direct path is available. If the wavelength convertor is full (i.e., it is occupied by the other lightpaths) $c = \infty$ is considered. When a connection request arrives, the algorithm calculates the cost for each possible path, and chooses the one with the lowest cost. A random selection is performed for breaking the ties [1].

Least-Congestion-Path. This method determines a list of paths for each source-destination pair. When a connection request arrives, the path with the lowest load amongst the corresponding pre-determined paths is chosen. In case of a tie, the other algorithms such as the shortest path routing algorithm may be performed. Since this algorithm considers all links on all pre-determined paths the computational complexity is high. For reducing the complexity of this method, the author of [6] proposed to only consider the first k links on each path (referred to as source neighborhood information) where k is defined appropriately.

2.2 Fault Tolerant Routing

A common approach to protect connections against link (node) failure is to consider at least two link-disjoint (node-disjoint) paths for each source-destination pair. In case of failure the backup path can be used. Fixed alternate routing directly provides some reservations for each connection. In the case of adaptive routing, backups should be considered by a protection scheme, which determines the alternate routes immediately after the primary connections have been established. The protection scheme can be the same as routing scheme with considering the cost of ∞ for undesirable links to promote link-disjoint (node-disjoint) paths [7].

2.3 Wavelength Assignment

Wavelength assignment is the second sub-problem within the RWA problem, and is usually addressed as an independent problem. The main objective of the wavelength assignment problem is to assign a wavelength to each connection in an efficient way such that no two lightpaths on a link share a common wavelength. Wavelength assignment can be resolved for either dynamic traffic or static traffic.

2.3.1 Static Wavelength Assignment

In static wavelength assignment the lightpaths and their routes are known in advance and we need to assign wavelengths to each lightpath, such that each of the lightpaths on a given link occupies a unique wavelength. This problem can be presented in two forms. For a given number of wavelengths the objective is to maximize the number of connections which can be established. For a given topology (without wavelength convertors) the main objective is to minimize the number of wavelengths needed for a set of connection requests under the wavelength continuity constraint. The latter problem can be reduced to a sequential graph coloring problem [8].

Graph Coloring Problem. The problem of static wavelength assignment can be reduced to a graph coloring problem which is NP-Complete; among the NP problems it is the least likely to be solved by any polynomial-time algorithm, because if there is any algorithm which solves this problem quickly, any problem in NP can be solved quickly by the same algorithm [4]. This proves that the wavelength assignment problem is NP-complete itself.

The graph coloring problem constructs a graph G of nodes, such that each node corresponds to a lightpath in the network. Each common fiber between two or more lightpaths in the actual system is represented by a link between their corresponding nodes in the graph. The next step is to color the nodes of the graph in such a way that no two adjacent nodes have the same color [8]. The minimum number of colors required by the process for coloring the graph G (the minimum number of wavelengths) is not easy to determine. Therefore heuristics such as the sequential graph-coloring algorithm have been introduced which are quite efficient in practice [9].

2.3.2 Dynamic Wavelength Assignment

Dynamic wavelength assignment considers the more realistic case where connection requests arrive dynamically. The connection requests are accepted if sufficient resources are available for the complete path and are blocked otherwise. Dynamic algorithms that consider a fixed number of wavelengths try to minimize the blocking probability. We now summarize some of the proposed heuristics, drawing from the review material in [1].

Random Wavelength Assignment. When a connection request arrives, a search procedure determines the available wavelengths on the appropriate path. One of the available wavelengths is randomly selected .

First-Fit(FF). The FF algorithm numbers all of the wavelengths in the network. When a connection request arrives, the first available wavelength is assigned to that connection. This approach has a very low computational complexity.

Most-Used(MU)/Pack. This method has been proposed for networks with fixed routes between each source destination pair [10]. The idea is to pack connections into a fewer number of wavelengths by selecting the most-used wavelength in the network. Packing connections is valuable in networks with the wavelength continuity constraint. In these networks there is no wavelength conversion and reserving the least-used wavelengths reduces the blocking probability in the network.

Least-Loaded(LL). LL has been designed as a wavelength assignment technique for multi-fiber networks [11]. For each connection request the most-loaded link of multi-fibers along the required path is looked over by the algorithm to determine the least-loaded wavelength on this link. The least-loaded wavelength on a multi-fiber network is the one with the largest residual capacity. The residual capacity of each wavelength may vary between 0 and N, where N is the number of fibers on each link.

MAX-SUM(MΣ). This method also has been proposed for multi-fiber networks [10, 12]. MΣ considers all the routes in the network that might be used by any of the connections. Then for a request the algorithm selects a wavelength on the selected path in such a way that the maximum available capacity on the remaining paths is obtained.

Wavelength Reservation(RSV). A wavelength reservation scheme reserves a specific wavelength on the links along a path of a multi-hop stream. This approach reduces the blocking probability for multi-hop streams, although the blocking probability for single-hop streams may increase.

3 Switching Schemes in WDM Systems

In addition to routing and wavelength assignment several different technologies have been developed for the transfer of optical data over WDM networks, such as optical packet switching, optical burst switching, and photonic slot routing. In this section we look at these switching technologies.

3.1 Photonic Slot Routing

Photonic Slot Routing(PSR) is a time division multiplexing approach for all-optical access and metro networks. PSR attempts to reduce complexity by eliminating the use of individual wavelength switching (IWS) components. The time-shared nature of this approach provides a sufficient level of traffic aggregation in networks for which the wavelength routing solution is inefficient. We now examine a photonic slot routed network which has been designed by Zang et al. [13].

3.1.1 Photonic Slot Routing in All-Optical WDM Mesh Networks

In PSR, time is slotted into fixed spans, each comprising a photonic slot. A photonic slot includes all wavelengths in a network. The packets of data destined for the same node are loaded into the photonic slots and are sent as a single integrated unit. Therefore there is no need for individual wavelength routing along a path and wavelength insensitive components are adequate for routing the photonic slots. Eliminating the use of wavelength demultiplexers results in faster switching, less complexity and lower cost [13].

Network Architecture. In this design a mesh network of wavelength insensitive nodes (Figure 1) is considered. At the source end, each node considers a separate electronic buffer space for each destination. The photonic slots for each destination consist of several data packets on a number of wavelengths and a header on a different wavelength. At each intermediate node headers are extracted from slots. During the header processing period the data slots travel along delay lines. To avoid the need for delay lines a solution is to send the header of a slot a fixed period of time before the payload, long enough for processing the header. The headers contain information about the wavelengths being used by the slot and the destination addresses. When a header of a slot arrives at a node which has some packets headed for the same destination, the node may insert its packets to the free wavelengths of the arriving slot. Inserting the packets can be done by using couplers.

When two or more arriving slots contend for the same output port several techniques may be used such as optical buffering, deflection routing, or dropping the slots randomly.

3.1.2 PSR Protocol

When a node has several packets to send, it may choose to add its packet to an arriving slot headed for the same destination or it may place the packets in an empty slot and sign the slot for its destination. There are a number of policies for slot assignment and adding packets to the existing slots. A slot may be occupied entirely by a node or may be left with some free spaces to be used by intermediate nodes. Here two types of slot assignment policies are introduced.

Packet Arrival Based Assignment Policy. Upon receiving an empty slot a node randomly chooses one of its queues and inserts a number of packets into the slot. This policy implies that as long as there are empty slots or free space in arriving slots with the appropriate destination the node is allowed to insert its packets. However this approach results in unfairness in resource allocation as well as high probability of blocking. Nodes located in the internal regions of the network usually receive assigned or full slots, while nodes located towards the edges of the network usually receive empty slots.

Capacity Allocation. A straightforward method for capacity allocation is the slot preassignment approach proposed by Chlamtac et al. [14]. In this approach a TDMA frame consisting of L slots is considered, in which a fixed number of slots is assigned to each source-destination pair. The number of slots for each pair is determined by using a network-wide TDMA schedule to achieve fairness and contention free slot routing at intermediate nodes. However this approach is not practical

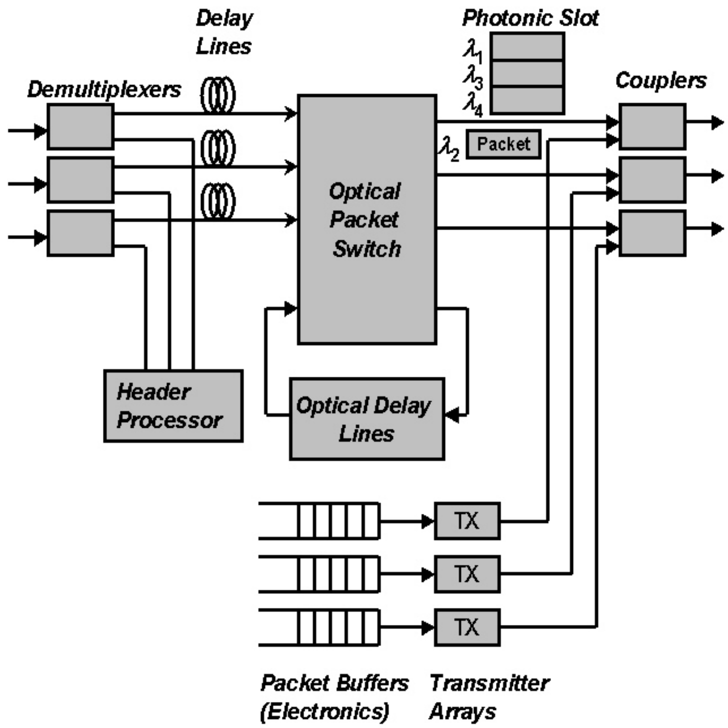

Figure 1: Photonic Slot Routing (PSR) node architecture; each node is capable of buffering, header processing, and packet insertion. A packet on wavelength λ_2 is inserted into an arriving photonic slot which is free at λ_2 [13].

since the length of the frame is fixed and the number of slots per pair is always an integer number. An alternative method [13] is to assign destinations probabilistically to arriving slots based on the capacity allocation results. This approach consists of two steps: the "Capacity-Allocation" step in which the fraction of the capacity of a link for each source-destination pair is determined, and the "Slot-Assignment" step in which a destination is assigned to each slot based on the results of the first step.

3.2 Optical Packet Switching

Optical packet switching (OPS) offers a better bandwidth granularity compared to the circuit switched networks, which results in a finer transmission, a bandwidth-efficient design and a more flexible all-optical network. However, this technology faces a number of limitations, in that optical packet switching requires optical buffering and packet-level processing [15].

An optical packet network consists of interconnected optical packet switches which are mainly composed of four parts: the input/output interfaces, switching fabric and control unit. The input interfaces align the packets, extract the header information from the packets and remove the headers. The control unit arranges the control functions based on the header information. The switch fabric switches the packets optically based on the control information. The output interfaces are responsible for optical signal regeneration and header insertion [15]. Header processing must be performed in the electrical domain, while the payload remains in the optical domain. Therefore packets have to be stored in the switch (e.g., using Optical Delay Lines (ODLs)) to be forwarded to the next stage when the header processing is complete, referred to as store-and-forward nature of packet switching.

3.2.1 Contention Resolution in Optical Packet Switching

Whenever two or more packets are destined to the same output port at the same wavelength, external blocking occurs. Optical packet switches are mostly non-blocking designs, therefore internal blocking is not an issue. To resolve contention in optical packet switches several approaches have been introduced including optical buffering, wavelength convertors, and deflection routing.

Optical Buffering. Optical delay lines are currently the only way of implementing optical buffers. Using electrical buffering is not acceptable, since electronic components can-not catch up with optical speeds. Optical buffers can be implemented by several ODLs with variable lengths to provide different delay lines. A counter is used to keep track of the number of packets in the buffer. Using a separate counter for each ODL adds flexibility to the use of optical buffers and reduces the length of the delay lines (i.e., a packet may be circulated through several ODLs to achieve the desired delay). Implementing optical buffers needs an enormous amount of fiber and a complex electronic control. In addition, optical signals travelling through delay lines may experience a considerable power loss so that optical amplifiers are necessary. The accumulated noise due to cascaded amplifiers limits the network size or requires signal regeneration, which is expensive. For more details about recent advances in optical buffering see Shun et al. [16] and Harai et al. [17].

Wavelength Convertors. Wavelength convertors can be used to reduce the number of delay lines, by converting the wavelength of a contending packet to a free wavelength at the output port. In a switch capable of optical buffering and wavelength conversion, the input is demultiplexed and each packet is sent to an appropriate destination (e.g., an output port). In the case of contention, a non-blocking space switch may send a packet on an available wavelength at the output port by using wavelength convertors or it may delay the packet. However wavelength convertors are expensive and full conversion is not easy to achieve [16].

Deflection Routing. Deflection routing resolves the contention by sending one of the contending packets to the desired link and passing the rest through any available link. The deflected packets are routed at the other nodes to their destination. This way packets of the same source-destination may experience different routes with different number of hops, which affects the network performance. Deflection routing usually is used in conjunction with optical buffering to reduce the need for buffering and to avoid too many recirculations at delay lines, which gives rise to signal to noise ratio degradation [16]. In the simplest method of deflection routing, delay lines are not used at all (see the hot-potato routing approach [18]).

3.2.2 Synchronous and Asynchronous Optical Packet Networks

Optical packet networks can be divided into two main categories: synchronous (slotted) and asynchronous (unslotted) frameworks. In the slotted case all the packets are aligned before they enter the switch. Whether or not the network is synchronous, bit-level synchronization and fast clock recovery are necessary at the switching stage for header recognition and packet delineation [16].

Synchronous Optical Packet Networks. In a slotted network the packet size is fixed. Each packet and a fixed guard time should fit into a fixed length time-slot. Therefore the packets arriving at the switch are aligned in phase with a local clock reference. In such a network contention is mostly resolved by using optical delay lines with propagation delays equal to a multiple of the time-slot duration [16].

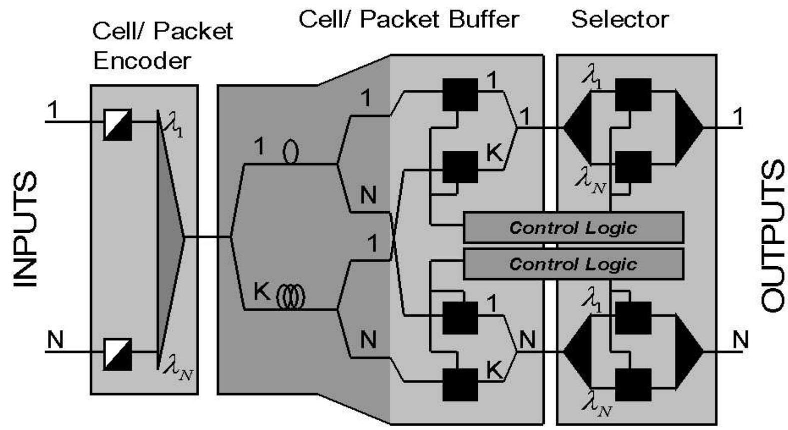

Figure 2: Schematic of Broadcast-and-Select Switch including a wavelength encoder, a cell/packet buffer and broadcast section, and a wavelength selector block (see section 3.2.3) [19].

Asynchronous Optical Packet Networks. In an unslotted network the packets are not aligned and switching the packets may be placed at any time. Also the packets are not required to have the same size. Therefore the general behavior of the network is more unpredictable and the chance of contention is higher than that of slotted networks. On the other hand such a network is easier and cheaper to implement, and the network is more flexible in case of control and contention resolution [16].

3.2.3 A Switch with KEOPS Architecture and Broadcast-and-select fabric

The KEOPS (KEys to Optical Packet Switching) switch was proposed as one of the key architectures for the European ACTS KEOPS project [19]. This $N \times N$ switch includes a wavelength encoder, a cell/packet buffer and broadcast section, and a wavelength selector block (Figure 2). At each input the wavelength convertor encodes its packet on a fixed wavelength. The cell/packet buffer block contains a splitter, K optical delay lines and a space switch stage. The space switch stage consists of a splitter, optical gates or Clamped-Gain Semiconductor Optical Amplifiers (CG-SOA) and combiners. CG-SOAs select the correct packets from delay lines and send them to their corresponding outputs based on the control unit information. The wavelength selector block is responsible for amplifying the signals before retransmission. Each block corresponding to each output selects the correct packets from inputs (i.e., the packets destined to that output port) and sends them to appropriate SOAs, which are wavelength dependent amplifiers. At the output ports the signals with different wavelengths are recombined and sent out. In this switch each input and output carries only one wavelength but the wavelength for an output port is not fixed (i.e., it can vary from packet to packet).

3.2.4 A Switch with a Broadcast-and-Select Fabric and Recirculation Buffer

A switch with a broadcast-and-select fabric proposed in [20] is shown in Figure 3. As in the architecture described in section 3.2.3 the wavelength at each output port depends on the incoming packets' wavelengths. In this switch fabric, at the node inputs, tunable wavelength convertors (TWCs) convert the wavelengths of the incoming packets to new wavelengths under the control unit information. At the beginning of each time-slot the inputs to the coupler are K ($K \leq N$)

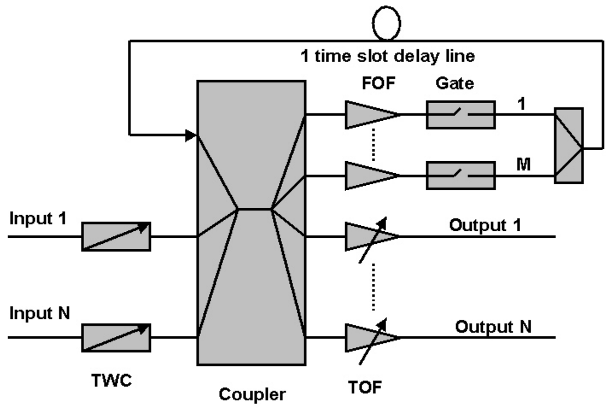

Figure 3: A Switch with a Broadcast-and-Select Fabric and Recirculation Buffer (see section 3.2.4) [19].

incoming packets plus $M-K$ inputs which are feedback from the 1 time-slot delay line. The coupler combines the input wavelengths and splits them into tunable optical filters (TOFs) ($N < M$ TOFs are available) and fixed optical filters (M fixed optical filters are available). According to the packets' destination the control unit determines up to N packets to be sent out through N TOFs and the remaining packets are delayed and recirculated through 1 time-slot delay line. The recirculated packets form part of the input to the coupler at the beginning of the following slot.

3.2.5 Architecture with a Wavelength Routing Switch Fabric

Many switch architectures based on wavelength routing have been proposed in the literature [21]. The switching procedure in most of them includes using ODLs for contention resolution, and routing the packets to the correct output ports through the wavelength switch fabric. Here we review an architecture with an input-buffered switch fabric [22]. As Figure 4 shows this switch is composed of scheduling and switching sections. In the scheduling section each incoming wavelength is passed through a TWC and then the two $K \times K$ arrayed waveguide gratings (AWGs), where $K = max(N, M)$. The AWGs are responsible for switching the optical signals from input ports to output ports. Each incoming and outgoing port carries a single wavelength. At each output port the wavelength varies with the packet. Between the two AWGs, M ODLs are used to resolve either internal or external contention. This architecture implements N individual buffers (corresponding to N different incoming wavelengths), each of which incorporates M positions. Each TWC converts the wavelength of the incoming packet to another wavelength in such a way to meet an ODL with an appropriate delay (each delay line has a different fixed wavelength). The ODLs are selected so that no two packets appear at the output of any ODL or the switch at the same slot. After the packets experience appropriate delays they are forwarded to their destination output ports through TWCs and AWGs. TWCs assign the appropriate wavelength to each outgoing packet corresponding to its destination port. Using wavelength convertors reduces complexity in the switching section.

3.3 Optical Burst Switching

Optical Burst Switching (OBS) tries to combine circuit and packet switching while avoiding the shortcomings of each. OBS is based on a one-way reservation scheme in which bursts of data follow control packets without waiting for acknowledgment [23]. Using OBS there is no need for optical-electrical-optical (O/E/O) conversion as in optical circuit switching, and similar to packet

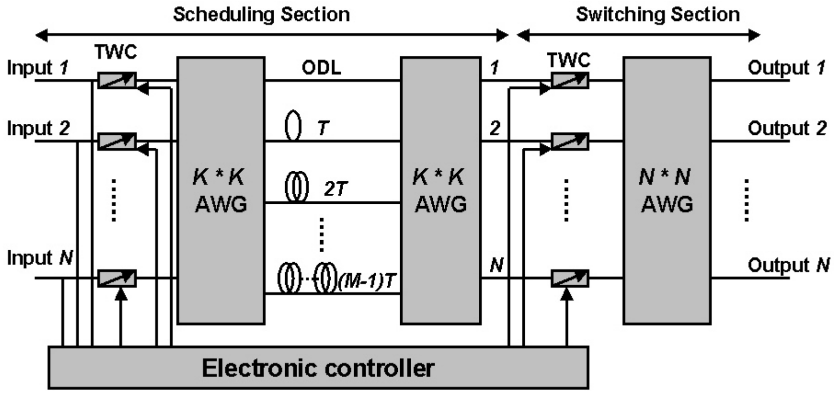

Figure 4: An input-buffered switch including Array Waveguide Gratings (AWGs) for switching the optical signals and Tunable Wavelength Convertors (TWCs) [22].

switching techniques it enables the network to share the resources among a number of users and increase the bandwidth utilization. In OBS the control packets are sent before the burst of data, so they are processed before the data arrives at the intermediate nodes. Consequently there is no need for data buffering as in packet switching. Moreover aggregating the packets as a burst reduces the overhead due to control packets. However, OBS is subject to a significant packet loss and burst retransmission. In the following we describe some proposed approaches for this switching method.

3.3.1 Optical Burst switching for Service Differentiation in the Next-Generation Optical Internet

The author of [23] surveys various designs for an optical burst switch network . In OBS the control packets representing a burst of data are processed at each node along a predetermined route to establish a lightpath by configuring the WDM switches. Then the corresponding burst passes through the pre-configured lightpath. There are three different techniques for burst switching, called, tell-and-go (TAG), in-band-terminator (IBT), and reserve-a-fixed-duration (RFD). Among these techniques the third one has been studied for all-optical networks [23]. RFD is the basis for an optical burst switching protocol called Just-Enough-Time (JET) [24], which is described as follows:

An OBS Protocol using Offset-Time and Delayed-Reservation (Just-Enough-Time). The basic functionality of this protocol is shown in Figure 5. The header of each burst is sent before the payload by a "base" offset time, $T \geq \Sigma_{h=1}^{H}\delta(h)$, where $\delta(h)$ is the expected control delay at hop $1 \leq h \leq H$. In Figure 5, H=3 and $\delta(h) = \delta$. In the JET-based control protocol the bandwidth at each hop is reserved from burst arrival time, t_s, until burst departure time, $t_s + l$, where l is the length of the burst. At hop i, the burst arrival time is the summation of t_a, the time at which the control packet processing has been finished, and $T(i) = T - \Sigma_{h=1}^{i}\delta(h)$ which is the remainder of the offset time at hop i (i.e. $t_s = t_a + T(i)$). If the reservation fails due to contention with other requests, the burst will be blocked. Blocking may be resolved by using ODLs. If ODLs are not available then the burst will be dropped [23].

An Offset-Time-Based QOS Scheme. This scheme uses an "extra" offset-time for supporting class isolation (or service differentiation) with or without ODLs [25]. Suppose that there are only two classes of traffic, class 0 and class 1, where class 1 has the highest priority. For the class 1

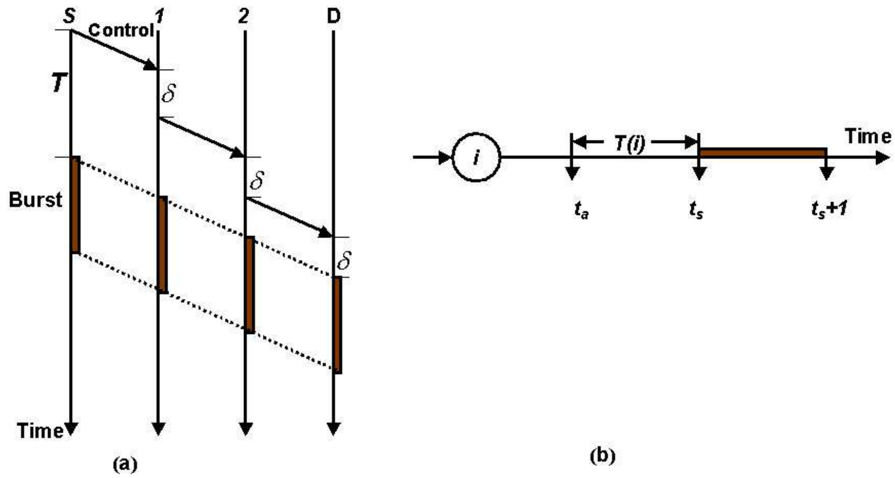

Figure 5: The use of offset time, and delayed reservation in JET-based OBS: (a) the burst follows the control packet after a base offset time, T; (b) the bandwidth is reserved from the burst arrival time, t_s [23].

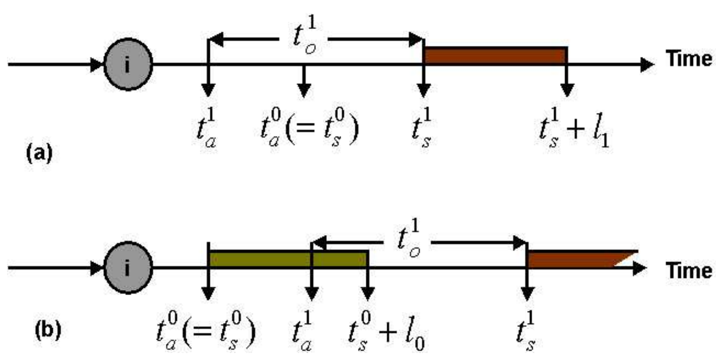

Figure 6: Class isolation at an optical switch without ODLs: (a) shows the case where request 1 has arrived before request 0; and (b) shows the case where request 0 has arrived before request 1 [23].

bursts, an extra offset time, t_o^1 is considered. The base offset time is considered negligible compared to this extra time. Let t_a^i and t_s^i be the arriving time and the service start time for a class i request, respectively, and l_i be the length of the request. To show how the class differentiation scheme works, we consider the following cases without ODLs (Figure 6):

Case 1, where the request 1 has arrived before request 0, $t_a^0 > t_a^1$ (Figure 6-(a)):
$if\ t_a^0 + l_0 \geq t_s^1 \geq t_a^0$ or $t_s^1 \leq t_a^0 \leq t_s^1 + l_1$ the class 0 request will be dropped
$else$, the request 0 will be succeed.

Case 2, where request 0 has arrived before the request 1, $t_a^1 > t_a^0$ (Figure 6-(b)):
$if\ t_a^1 + t_o^1 < t_a^0 + l_0$ request 1 will be dropped.

To avoid blocking of class 1 requests, t_o^1 needs to be larger than the maximum burst length in class 0. When ODLs are available the problem is more complicated but by choosing an appropriate offset time, class 1 can be isolated from class 0 requests in reserving both bandwidth and ODLs [23].

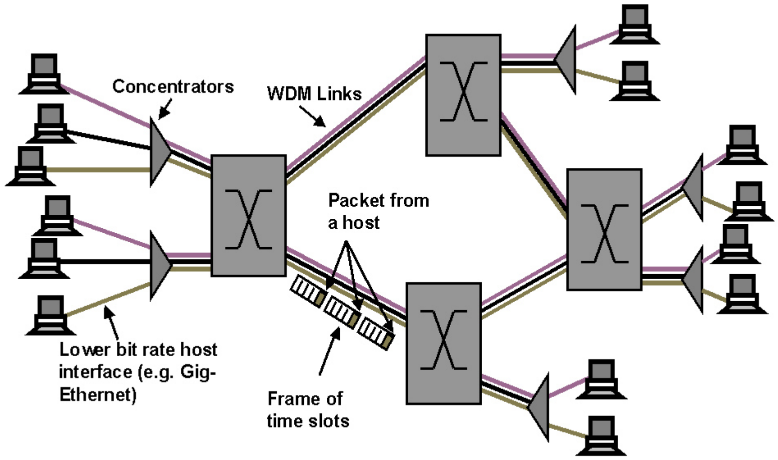

Figure 7: Time-sliced packet switched network architecture, composed of a network of TSOBSs and concentrators; concentrators transmit user data bursts in time-division channels. Data is carried on the WDM link within frames comprised of fixed length time-slots [7].

3.3.2 Time Sliced Optical Burst Switching

Time Sliced Optical Burst Switching (TSOBS) is an approach proposed by Ramamirtham et al. [7], which replaces wavelength domain switching with time domain switching. This eliminates the need for wavelength converters and results in a lower cost. TSOBS also reduces the number of delay lines required for supporting burst switching by using blocking Optical Time-Slot Interchangers (OTSI) and Optical Crossbars.

TSOBS Networks. Figure 7 shows the architecture of a time-sliced optical burst switched network. Terminals or other networks are connected to the TSOBS network by concentrators. WDM links connect the network of time-sliced optical burst switches. The information on each wavelength is organized into a series of frames of fixed length time-slots. Concentrators transmit user data packets (packet switching) or aggregated user data (burst switching) in time-division channels. The Burst Header Cells (BHC) are carried on separate control wavelengths. If the ratio of the average burst length to the BHC length is L, each control wavelength can support $L-1$ data wavelengths. Each BHC carries information about the length of the burst, source-destination addresses, wavelength, and the identification of the first frame in which the data burst appears. It also includes a field in which the distance travelled by a burst and the number of optical operations since its last regeneration are recorded. This information is used to regenerate the burst before too much signal degradation arises due to noise and signal attenuation at each switching stage [7].

Optical time-slot Interchangers. OTSIs perform time domain switching in the TSOBS architecture. They can be used in either non-blocking or blocking designs. The simplest form of a non-blocking OTSI has N delay lines and an $(N+1)\times(N+1)$ optical crossbar switch (see Figure 8) , where N is the number of time-slots in each frame. Delay lines provide $1, 2, \cdots, N$ time-slots interval delay. A more practical design for a non-blocking switch has been introduced in [7]; the design uses $2\sqrt{N}-2$ delay lines and a $(2\sqrt{N}-1)\times(2\sqrt{N}-1)$ crossbar switch. In order to reduce the cost and complexity of switches, blocking OTSIs can be used. For this design $N/2$ delay lines provide $1, 2, 4, \cdots, N/2$ time-slot delays. A small crossbar switch $(\sqrt{N}\times\sqrt{N})$ is needed to support switching operations between the delay lines as well as the incoming time-slots. Thus the number of

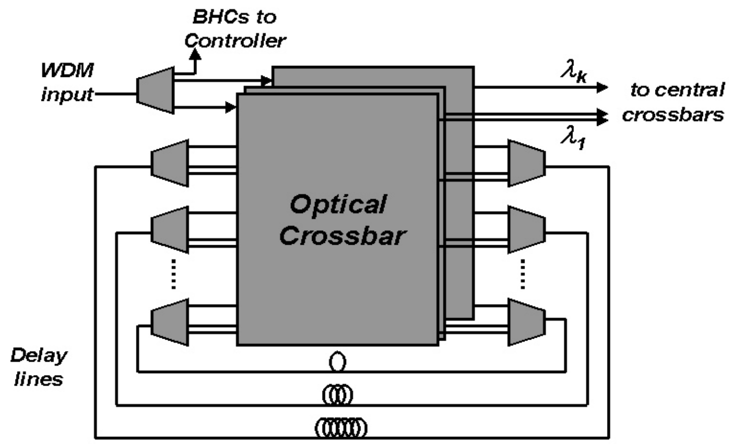

Figure 8: Optical time-slot Interchanger, which performs time-domain switching; each OTSI contains a set of optical crossbars for switching time-slots among the inputs, outputs and a set of delay lines. The signals are demultiplexed to perform the switching operations and remultiplexed onto the delay lines, allowing the cost of the delay lines to be shared by the different wavelengths [7].

switching operations increases as the switch is subject to blocking. When a time-slot arrives at the switch, a search procedure is performed to find an available sequence of delay lines for creating the desired delay for that slot. For this task the state of the delay lines must be tracked and recorded in a scheduling array. To find the best delay path, a shortest path tree algorithm is performed. This algorithm finds the path with the smallest number of switching operations which is needed to obtain the desired delay. Blocking OTSIs can be presented with different number of delay lines and crossbar switches. Therefore each design has a different cost and blocking probability [7].

Switch Architecture. Figure 9 shows the overall design for a TSOBS router. The router consists of synchronizers, OTSIs, a central optical crossbar, a controller and WDM multiplexors. The synchronizers can be implemented using a space-division optical switch and a finely calibrated set of delay lines to synchronize the incoming frames to the local clock. The performance evaluations show that the blocking design performs very close to non-blocking switches while the cost is improved significantly due to a small number of delay lines needed to provide an acceptable blocking probability. OTSI separates the control wavelengths and forwards them to the controller. The controller processes the BHC information and passes the results to optical crossbar switch as well as the OTSIs. When a BHC arrives at the controller, the controller determines the appropriate outgoing link for that burst. Then on the outgoing links and the output of the corresponding OTSI a look up procedure defines the available time-slots on the wavelength being used by the burst. OTSI then provides time domain switching for all wavelengths available on each link based on the information about the available time-slots on the outgoing links. It separates the data wavelengths and sends each one on a different fiber to the optical crossbar after the appropriate time domain switching is performed. The number of delay lines used in OTSIs is the key parameter affecting both the performance and the cost of the TSOBS switches. The optical crossbar performs space switching on individual wavelengths and WDM multiplexors combine the data wavelengths with control wavelengths on the outgoing links.

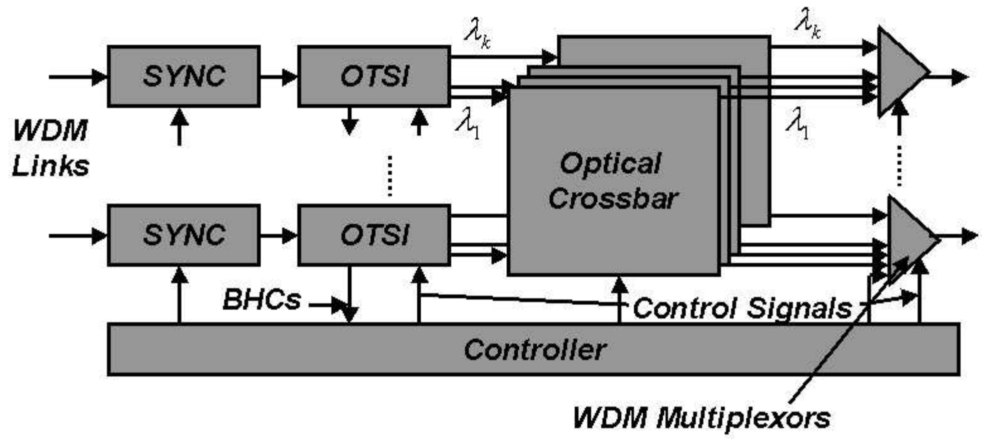

Figure 9: The overall Time-Sliced Optical Burst Switch Design; each incoming WDM link terminates on a Synchronizer (SYNC) which synchronizes the incoming frame boundaries to the local timing reference. Optical Time Slot Interchangers (OTSI) provide the required time domain timing of switching operations on the data wavelengths [7].

4 Bandwidth Reservation in WDM systems

In an all-optical communication system, wavelength division multiplexing (WDM) has been introduced to use the substantial optical bandwidth efficiently by partitioning the optical bandwidth into WDM channels compatible with the speed of electronic processing. WDM systems can incorporate time-division multiplexing (TDM), which provides WDM access to many end-users sharing the optical channels in the time domain, and therefore makes optical networking practical.

In order to effectively facilitate multiple access to the network resources, bandwidth reservation can be used as an admission control mechanism. Bandwidth reservation uses at least one channel for reservation and the rest are allocated to data transmissions [26]. Reservation can be performed in several forms including burst switching and fixed time-slot assignment. In a burst switching model the bandwidth is requested before transmitting the burst of data, while the data is usually transmitted after a period of time, without waiting for acknowledgment [23]. The waiting time is considered just long enough for processing the request of bandwidth on an appropriate outgoing link, which should be done in relation to the other contending requests. In fixed time-slot scheduling the time is divided into small partitions called slots, which are assigned to the requests ahead of time. Therefore in this model the transmission is performed after receiving the approval from the controller. In time-slot scheduling, the assignment can be accomplished using many varieties of scheduling algorithms.

Scheduling algorithms can be designed in an *on-line* [26–28] or *off-line* [29–31] manner, each of which involves a different set of requirements and techniques. An on-line (or incremental) scheduling algorithm computes a schedule based on the available partial information for each arriving request. During the reservation phase, as soon as the first request arrives the scheduler starts computation. This technique has a relatively low computation time, but the result is not always optimal [26]. On the other hand the off-line scheduling algorithms need the entire traffic demand matrix to start computing the schedule. The demand matrix whose (i, j)-th element represents the number of (fixed-size) packets that must be transmitted from source i to destination j, can be collected during the reservation phase (in a dynamic traffic scenario) or the algorithm might consider an *a priori* traffic distribution pattern (in a static traffic scenario) [27]. In the latter case referred to as *static* scheduling, the traffic demand changes do not affect the schedule [26].

Depending on the traffic pattern the scheduling problem can be solved for a fixed frame length or a variable frame length. The variable frame length structure does not suit connections which need reservations for several consecutive frames. In long term transmissions the users usually specify their requests in bits per second. Therefore changing the frame length changes the user's throughput if the number of slots per frame allocated to a user is not changed. Consequently the variable frame length formalization is more suitable for one-shot transmissions [32].

4.1 Bandwidth Reservation in Star Coupled Networks with tunable transmitters/receivers

Many algorithms have been proposed so far for bandwidth reservation in an all-optical broadcast-and-select network with a star topology [30, 32–36]. In this network a star coupler provides the connections between several nodes, which are equipped with tunable transmitters and/or tunable receivers. In general, the number of nodes is greater than the number of available channels, and therefore the transmitters have to share the channels using a reservation-based technique. Depending on the efficiency requirements and the traffic pattern the scheduling problem can be solved for either off-line or on-line formalization and variable or fixed frame length.

4.2 Variable Frame length

The scheduling problem for WDM networks with tunable transmitters /receivers and variable frame length has been extensively analyzed [30, 32–34]. The proposed algorithms try to optimize the schedule while delivering all traffic within the next frame. This problem is usually formalized as follows [35]:

> "Given a traffic matrix D whose elements D_{ij} are the numbers of (fixed-size) packets that must be transmitted from any source user i to any destination user j, find a time/wavelength assignment that guarantees the delivery of all traffic, while minimizing the time necessary to accommodate all transmissions (the frame duration), subject to tuning delay constraints."

The main objective of a scheduling algorithm is to minimize the computation time while maximizing the utilization of the network resources. In this particular network with a variable frame length, increasing the utilization is equivalent to reducing the schedule length. The lower bound on the schedule length is given by $\max_{i,j}\{S_i, R_j\}$, where $S_i = \sum_{j=1}^{M} D_{ij}$ is the i-th column sum and $R_j = \sum_{i=1}^{M} D_{ij}$ represents the j-th row sum. It has been shown that the algorithms aimed at minimizing the transmission schedule length have polynomial-time complexity [37]. Also it has been shown that minimizing the effect of tuning latency on the schedule length while minimizing the transmission time is an NP-hard[4] problem [38]. In [30, 32–34] several heuristic approaches have been proposed which aim at minimizing the frame duration with the assumption that all of the requests are assigned. In the following we introduce two simple algorithms for the variable-frame problem in a star-coupled network with tunable transmitters and receivers.

4.2.1 Heuristic Approaches

We now review two reservation-based scheduling techniques for WDM star networks [39]: SEQSAM (SEQuential Scheduling AlgorithM) and BALSAM (BALanced Scheduling AlgorithM). The net-

[4] A problem is NP-hard if solving it in polynomial time would make it possible to solve all problems in class NP in polynomial time.

work architecture is based on a passive star topology composed of M nodes which are equipped with tunable transmitters and receivers capable of operating on C channels. Frame transmission includes a reservation phase, a schedule computation phase, and a data phase. During the TDM-based reservation phase, all the nodes broadcast their information (i.e. a control packet containing the destination identification and the requested packet size is sent to every node in the network). This can be done by tuning all the transmitters and receivers on a single channel. At the end of each reservation phase an $M \times M$ demand matrix, D, whose (i, j)-th element, D_{ij} denotes the demand from node i to destination j, is available. During the schedule computation phase a transmission schedule is computed at every node. The scheduling problem is to specify the durations in which each transmitter and receiver should tune to a specific channel during the data phase.

SEQSAM allows receivers to tune on multiple channels during each transmission phase. But this technique has a very low performance and is introduced in [34] primarily to show the performance improvement that can be achieved by the other algorithm, BALSAM. BALSAM restricts the receivers to tune on only one channel during each transmission phase. This restriction reduces the effect of tuning latency on the transmission time.

SEQuential Scheduling AlgorithM (SEQSAM). SEQSAM groups the elements of the $M \times M$ demand matrix into groups of C elements, producing G sub-matrices. Each of the sub-matrices has at most C nonzero elements with no more than one nonzero element on any row or column. Therefore the lower bound on G is given by $(M^2 - M)/C$. SEQSAM implements a simple technique for obtaining the matrix decomposition, but its average time complexity is $O(M^3)$. Scheduling in SEQSAM takes place after grouping the elements and obtaining G sub-matrices:

$$D = D^1 + D^2 + D^3 + \cdots + D^G$$

The length of the transmission phase is the sum of the largest entries in the sub-matrices (D^i_{max}) defined as $\sum_{i=1}^{G} D^i_{max}$.

Example 1. Consider a star network with $M = 4$, $C = 2$, and the following demand matrix and its decomposition:

$$\begin{aligned} D &= \begin{bmatrix} 0 & 3 & 2 & 2 \\ 1 & 0 & 4 & 1 \\ 3 & 2 & 0 & 1 \\ 1 & 1 & 3 & 0 \end{bmatrix} \\ &= \begin{bmatrix} 0 & 3 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 2 \\ 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \\ &+ \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 3 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 3 & 0 \end{bmatrix} \end{aligned}$$

In this example the matrix is decomposed into $G = 6$ groups of $C = 2$ nonzero elements. Each sub-matrix is assigned a section of time slots on the transmission schedule. The length of each section is equal to the largest entry of the corresponding sub-matrix. On each sub-matrix the first channel is assigned to the first non-zero entry. The next channel is assigned to the next non-zero

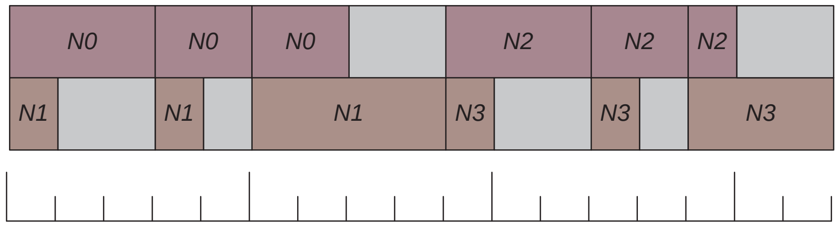

Figure 10: Example 1: Allocation schedule of two available wavelengths (shown by two different colors) achieved by SEQSAM. Ni denotes the source *node* i, and the schedule length is 17 time slots.

entry and so on. The schedule length (as we see in Figure 10) is 17 slots.

BALanced Scheduling AlgorithM (BALSAM). This algorithm converts the $M \times M$ demand matrix to an $M \times C$ matrix by using the Modified Multi-FiT (MMFT) algorithm [40]. Then using an interval-based scheduling algorithm it assigns the channels to the transmitters. BALSAM attempts to reduce the schedule length by reducing the effect of tuning latency and the wasted resources during each frame. It has been shown in [36] that the time complexity of the Interval-Based Scheduling Algorithm (IBS) is $O(MC^2K')$ where K' is the largest element in the $M \times C$ demand matrix. A simple description of this algorithm is as follows. At the first step the MMFT algorithm computes the column sums of the demand matrix and sorts the columns in descending order. Then it assigns each of the first C columns to each of the C channels and starts over with assigning the second C columns to the C channels and so on. So for $C < M$ several columns are given a common channel, and each receiver (corresponding to each column) has to tune on the selected channel. With this technique the load is almost equally distributed on the available channels, and we have an $M \times C$ demand matrix, whose (i, j)-th element represents the number of requested time slots from *node* i on *channel* j. The IBS algorithm for the $M \times C$ demand matrix keeps track of the available intervals on the channels. When a node request for a channel is considered, the algorithm tries to fit the request in the first available interval. If the interval is not sufficient to be allotted to the demand the next available interval is considered. The following example explains how this algorithm operates.

Example 2. Consider the demand matrix in the previous example. Using MMFT the 4×4 demand matrix is converted to the following 4×2 demand matrix which shows the demand of each transmitter for each channel. For more information regarding the MMFT algorithm refer to [39].

$$\begin{bmatrix} 4 & 3 \\ 5 & 1 \\ 1 & 5 \\ 3 & 2 \end{bmatrix}$$

Node 0 request: IBS first considers the request of *node* 1 for *channel* 1 which is 4 slots. Initially all of the slots are available. Therefore the first 4 slots are allotted to the first request of *node* 0. The next request to be considered is the *node* 0 request for *channel* 2. Since the transmitter has been tuned to *channel* 1 during the first 4 time slots, the first possibility for transmission on *channel* 2 is the interval [5,7].
Node 1 request: For the next request, *node* 1 on *channel* 1, the available intervals are $[5, \infty]$. Therefore the interval $[5, 9]$ is assigned to this node. Similarly the *node* 1 request for *channel* 2 has to

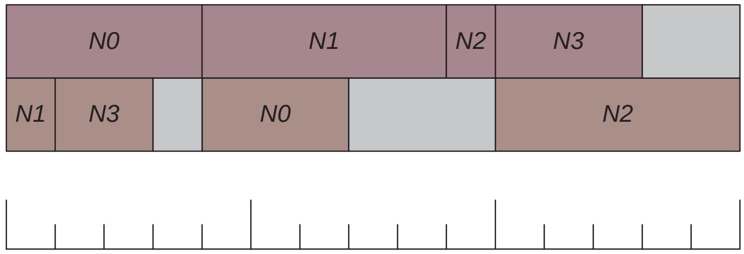

Figure 11: Example 2: Allocation schedule of two available wavelengths (shown by two different colors) achieved by BALSAM. Ni denotes the source *node* i, and the schedule length is 15 time slots.

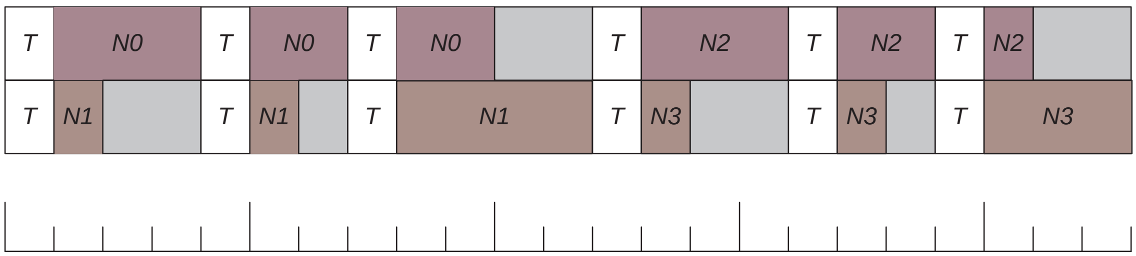

Figure 12: Example 1: Allocation schedule of two available wavelengths (shown by two different colours/shadings) achieved by SEQSAM incorporating the tuning latency for T=1. Ni denotes source *node* i, and the schedule length is 23 time slots.

be assigned based on a look up procedure which determines the first fit interval in the scheduling table. In this example the first slot on *channel* 2 suits this request. For the rest a similar procedure is performed. Figure 11 shows the final transmission schedule. The schedule length in this example is 15 slots.

Tuning Latency. This section considers the effect of tuning latency on the transmission time for the two proposed algorithms. Let T_t denote the number of slots required for tuning the transmitter between the channels, and T_r denote the tuning latency of the receivers. We define $T = \max\{T_t, T_r\}$. For SEQSAM there are G sub-matrices and transceivers tune to their assigned channels between the sub-matrices. Therefore the additional delay introduced by SEQSAM due to tuning latency is GT (Figure 12).

BALSAM introduces less tuning latency to the transmission phase since the receivers assignment to the channels is static during each frame. The receivers can tune to their assigned channels as soon as the MMFT algorithm is performed to balance the requests. Moreover the transmitters may be able to tune their channels ahead of time, during the intervals the other transmitters are sending data. Figure 13 shows the effect of tuning latency on the schedule length in the BALSAM algorithm. The schedule length is 17 slots in this case.

4.3 Fixed Frame Length

In the networks comprising the frames with a fixed length the aim is to minimize the number of packets that can not be transmitted in the scheduled frame. Formalization of the off-line scheduling problem considering a fixed frame length is as follows [35]:

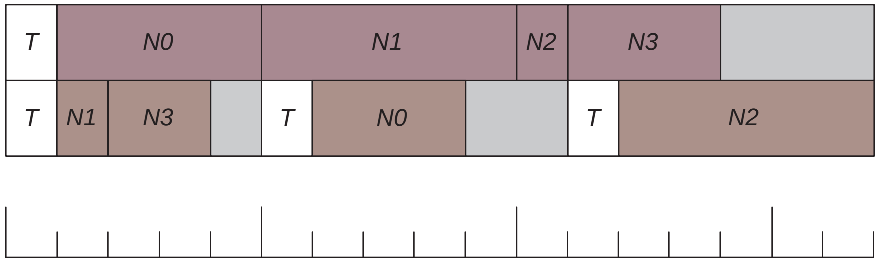

Figure 13: Example 2: Allocation schedule of two available wavelengths (shown by two different colours/shadings) achieved by BALSAM incorporating the tuning latency for T=1. Ni denotes source *node* i, and the schedule length is 17 time slots.

> "Given a traffic matrix D whose elements D_{ij} specify the number of packets that must be transmitted from any source i to any destination j in a pre-specified time frame comprising F slots, find a time/wavelength assignment (satisfying the tuning delay constraints), that minimizes the number of packets that are not accommodated in the frame."

This formalization states that a new scheduling is obtained by re-allocating network resources for all end-to-end user traffic flows, in other words, any change in bandwidth requests results in rescheduling all of the connections.

Even though off-line algorithms can lead to optimal solutions, the time necessary for receiving the whole traffic demand before starting the calculation is high. In order to reduce the time between two consecutive transmission phases, on-line scheduling algorithms have been introduced in which the computation phase is overlapped with the reservation phase by starting calculation as soon as the first request arrives. Then the requests arriving later are assigned the free time-slots without re-allocating the already allocated requests for the current frame.

In order to reduce the complexity and the bandwidth devoted to signalling, the on-line algorithms can accommodate a *transparency constraint* in the scheduling problem: new requests may be accepted only if they do not affect existing allocations, otherwise they are refused. Therefore the on-line scheduling problem imposed for a fixed frame with the transparency constraint is formulated as follows [35]:

> "Given a time frame comprising F slots, in which a number of user-to-user transmission are allocated according to a known schedule, and given a matrix D_n of new requests or modifications of allocated requests, find a time/wavelength assignment, (satisfying the tuning delay constraints), that avoids modification in existing allocations (except for those resulting from D_n) and minimizes the number of packets of D_n that are not accommodated in the frame."

In comparison, the off-line algorithms for dynamic traffic are more efficient in terms of the schedule they generate, but the computation time is high. The algorithms proposed in [30, 32, 33] compute static WDM/TDM schedules, which allocate the resources according to long-term bandwidth requirements. The algorithms are simple, but inefficient in the case of bursty traffic [35]. On-line algorithms usually have less computation time and lower efficiency, since the incremental nature of the algorithms does not provide in general an optimal schedule. In [35] Marsan et al. introduce an on-line algorithm which provides a tradeoff between simplicity and efficiency, assuming a slowly-varying traffic pattern. In the following section we review the scheduling algorithms proposed in [35].

4.3.1 Simple On-line Scheduling Algorithms for All-Optical Broadcast-and-Select Networks

Marsan et al. consider all-optical broadcast-and-select networks with a star topology which provides a number of slotted WDM channels for packet transmission [35]. Each node in the network includes a tunable transmitter and one fixed receiver. The range of the transmitters' tunability is sufficient for a full connectivity between each source/destination pair. A centralized controller provides time-slot assignment in a WDM/TDM frame considering long term bandwidth requests demanded by the users. Different strategies for on-line scheduling are proposed. The algorithms are executed periodically to re-compute the schedule in response to a change in the users' demands (i.e. a new request or a modification of an existing connection). When transparency is enforced, the existing connections in consecutive frames receive their previously assigned time-slots, and the new requests or the modified requests may occupy the free time-slots. The algorithms attempt to allocate the requests for each connection in contiguous slots in order to reduce the overhead due to tuning latency.

The slot allocation process for a new request of k slots to connection (i, j) starts by determining the wavelength on which a connection should be established. In the case of tunable transmitters and fixed receivers, the wavelength for connection (i, j) can be identified from the destination address. Then the algorithm searches for (i, j)-*eligible* slots on the destination wavelength, w_j. The (i, j)-eligible slots are defined as the set of free time-slots on wavelength w_j during which a transmitter is neither tuning nor transmitting on some other wavelength. The three different strategies presented in [35] use different criteria for selecting k slots among the (i, j)-eligible slots. In the first step the algorithms try to assign k consecutive slots to the request based on different criteria. If that is not possible, a second sequential search assigns k eligible slots to the demand on a first-fit basis, even though it is possible to apply a more sophisticated rule for the case that the request must be split. In the following we describe the different criteria for assigning the requests in the first step.

Algorithms Description:

Sequential Search (SS). The first strategy searches for the first k contiguous (i, j)-eligible time-slots, and assigns them to the request.The complexity of this algorithm is linear in the frame size F.

Best Fit Search (BFS). The second strategy searches for all of the sequences consisting of at least k contiguous (i, j)-eligible slots. Among these the shortest sequence is chosen. Similar to the first algorithm, the complexity of this algorithm is linear in the frame size, but it needs more memory for storing the information for selecting the shortest sequence.

Minimum Cost Search (MCS). The third strategy defines a reward function on the space of all free slots for each wavelength. For a demand of k slots on wavelength w_j a search procedure determines the set of k consecutive slots for which the global reward function after this assignment is maximum.

Denote by $C_{t,w_j} = N_{fs}(t) + N_{fw}(t)$ the reward associated with the (i, j)-eligible slot t on wavelength w_j, where N_{fs} is the number of free sources, and N_{fw} is the number of free wavelengths at slot t over all available wavelengths. Let F_{ij} be the number of (i, j)-eligible slots (on wavelength w_j), and S be the set of all contiguous source-free slots, the sequences on which source i is neither tuning nor transmitting. The reward associated with each sequence $s \in S$ of contiguous source-free slots is denoted by $\chi(s)$, an increasing function with s (e.g., $\chi(s) = 1.5\|s\|^{1.2}$, where $\|s\|$ indicates the number of slots in sequence s). Let W be the set of all free sequences on wavelength w_j, denoted as wavelength-free slots at this wavelength. Denote by $\psi(f)$ the reward associated with each sequence $f \in W$ of contiguous wavelength-free slots. The global reward function associated with source i and wavelength w_j is defined as:

$$M = \sum_{t \in F_{ij}} C_{t,w_j} + \sum_{s \in S} \chi(s) + \sum_{f \in W} \Psi(f), \tag{1}$$

When a new request must be assigned on wavelength w_j, a sequence of eligible slots is temporarily allocated to the request and M is computed. After computing M for all possible allocations, the MCS algorithm chooses the allocation which provides the maximum value of M. In other words this algorithm implements a search procedure which allocates a contiguous sequence of slots with the lowest impact on the value of the global reward function after the allocation.

References

[1] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *Optical Network Mag.*, vol. 1, pp. 47–60, Jan. 2000.

[2] G. Wedzinga, *Photonic Slot Routing in Optical Transport Networks*, Kluwer Academic Publishers, Boston, MA, Nov. 2002.

[3] G.V. Bochmann, T. Hall, O. Yang, M.J. Coates, L. Mason, and R. Vickers, "The agile all-photonic network: An architectural outline," in *Proc. 22nd Biennial Symposium on Communications*, Kingston, Canada, June 2004.

[4] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Co., June 1979.

[5] I. Chlamtac, A. Ganz, and G. Karmi, "Light path communications: An approach to high-bandwidth optical WANs," *IEEE Trans. Communications*, vol. 40, pp. 1171–1182, July 1992.

[6] L. Li and A.K. Somani, "Dynamic wavelength routing using congestion and neighborhood information," *IEEE/ACM Trans. Networking*, vol. 7, pp. 779–786, Oct. 1999.

[7] J. Ramammirtham and J. Turner, "Time sliced optical burst switching," in *Proc. IEEE Infocom*, San Francisco, CA, April 2003.

[8] B.Mukherjee, *Optical Communication Networks*, McGraw-Hill, New York, 1997.

[9] M. T. Jones and P. E. Plassmann, "A parallel graph coloring heuristic," *SIAM J. Scientific Computing*, vol. 14, pp. 654–669, May 1993.

[10] S. Subramaniam and R. A. Barry, "Wavelength assignment in fixed routing WDM networks," in *Proc. IEEE ICC*, Montreal, Canada, June 1997, vol. 1, pp. 406–410.

[11] E. Karasan and E. Ayanoglu, "Effect of wavelength routing and selection algorithms on wavelength conversion gain in WDM optical networks," *IEEE/ACM Trans. Networking*, vol. 6, pp. 186–196, April 1998.

[12] R. A. Barry and S. Subramaniam, "The max-sum wavelength assignment algorithm for WDM ring networks," in *Proc. OFC*, Dallas, TX, Feb. 1997.

[13] H.Zang, J. P. Jue, and B. Mukherjee, "Photonic slot routing in all-optical WDM mesh networks," in *Proc. IEEE Global Telecommunications Conf.*, Rio de Jan., Brazil, Dec. 1999, vol. 40, pp. 1449–1453.

[14] I. Chlamtac, A. Farago, and T. Zhang, "Lightpath (wavelength) routing in large WDM networks," *IEEE J. Sel. Areas Communications*, vol. 14, pp. 909–913, June 1996.

[15] L. Xu and and G. Rouskas H.G. Perros, "Techniques for optical packet switching and optical burst switching," *IEEE Communications Mag.*, vol. 39, pp. 136–142, Jan. 2001.

[16] Y. Shun, B. Mukherjee, and S. Dixit, "Advances in photonic packet switching: an overview," *IEEE Communications Mag.*, vol. 38, pp. 84–94, Feb. 2000.

[17] H. Harai, N. Wada, F. Kubota, and W. Chujo, "Contention resolution using multi-stage fiber delay line buffer in a photonic packet switch," in *Proc. IEEE Int. Conf. Communications*, April 2002, vol. 5, pp. 2843–2847.

[18] A. S. Acampora and I. A. Shah, "Multihop lightwave networks: A comparison of store-and-forward and hot-potato routing," *IEEE Trans. Communications*, vol. 40, pp. 1082–1090, June 1992.

[19] C. Guillmot, M. Renaud, and P. Gambini, "Transparent optical packet switching: The European acts keops project approach," *Lightwave Technology*, vol. 16, pp. 729–375, May 1998.

[20] G. Bendelli, M. Burzio, P. Gambini, and M. Puleo, "Performance assessment of a photonic ATM switch based on a wavelength controlled fiber loop buffer," in *Proc. Conf. Optical Fiber Communications (OFC)*, California, USA, Feb. 1996, pp. 106–107.

[21] J. M. Gabriagues and J. B. Jacob, "Photonic ATM switching matrix based on wavelength routing," in *Proc. SPIE Photonic Switching*, Minsk, Belarus, July 1992, vol. 1807, pp. 355–359.

[22] W. D. Zhong and R. S. Tucker, "Wavelength routing based photonic packet buffers and their applications in photonic packet switching systems," *Lightwave Technology*, vol. 16, pp. 1737–1745, Oct. 1998.

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

[24] C. Qiao and M. Yoo, "Just-enougth-time(JET): A high speed protocol for bursty traffic in optical networks," in *Proc. IEEE/LEOS Conf. Technologies for Global Information Infrastructure*, Montreal, CA, August 1997, pp. 26–27.

[25] M. Yoo, C. Qiao, and Sudhir Dixit, "Qos performance of optical burst switching in ip-over-wdm networks," *IEEE J. Select. Areas Commun.*, vol. 18, pp. 2062–2070, Oct 2000.

[26] K.M. Sivalingam, J. Wang, X. Wu, and M. Mishra, "An interval-based scheduling algorithm for optical WDM star networks," *J. Photonic Network Communications*, vol. 4, pp. 73–87, Jan. 2002.

[27] A. Bianco, M. Guido, and E. Leonardi, "Incremental scheduling algorithms for WDM/TDM networks with arbitrary tuning latencies," *IEEE Trans. Communications*, vol. 51, pp. 464–475, Mar. 2003.

[28] E. Johnson, M. Mishra, and K. Sivalingam, "Scheduling in optical WDM networks using hidden Markov chain based traffic prediction," *J. Photonic Network Comms.*, vol. 3, pp. 271–286, July 2001.

[29] W. Schmidt, "An on-board switched multiple-access system for millimetre-wave satellites," in *Proc. Digital Satellite Comm. Conf.*, London, UK, 1969.

[30] M. S. Borella and B. Mukherjee, "Efficient scheduling of nonuniform packet traffic in a WDM/TDM local lightwave network with arbitrary transceiver tuning latencies," *IEEE J. Sel. Areas. Comms.*, vol. 14, pp. 923–934, Sept. 1996.

[31] M. Azizoglu, R. A. Barry, and A. Mokhtar, "Impact of tuning delay on the performance of bandwidth limited optical broadcast networks with uniform traffic," *IEEE J. Sel. Areas Comms.*, vol. 14, pp. 935–944, Jun. 1996.

[32] M. Ajmone Marsan, A. Bianco, E. Leonardi, F. Neri, and A. Nucci, "Efficient multi-hop scheduling algorithms for all optical WDM broadcast-and-select networks with arbitrary transceiver tuning latencies," in *Proc. IEEE Globecom*, Sydney, Australia, Nov. 1998.

[33] A. Ganz and Y. Gao, "A time-wavelength assignment algorithm for a WDM star network," in *Proc. IEEE Infocom*, Florence, Italy, 1992.

[34] G. N. Rouskas and V. Sivaraman, "Packet scheduling in broadcast WDM networks with arbitrary transceiver tuning latencies," *IEEE/ACM Trans. Networking*, vol. 5, pp. 359–370, June 1997.

[35] M.A. Marsan, A. Bianco, E. Leonardi, F. Neri, and A. Nucci, "Simple on-line scheduling algorithms for all-optical broadcast-and-select networks," *IEEE European Trans. Telecommunications*, vol. 11, pp. 109–116, Jan. 2000.

[36] K.M. Sivalingam, J. Wang, X. Wu, and M. Mishra, "Improved on-line scheduling algorithm for optical WDM netwoks," in *Proc. DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, New Brunswick, NJ, 1998, pp. 43–61.

[37] Gopal I. S, G. Bongiovanni, M. A. Bonuccelli, D. T. Tang, and C. K. Wang, "An optimal switching algorithm for multibeam satellite systems with variable bandwidth beams," *IEEE Trans. Communications*, vol. 30, pp. 2475–2481, Nov. 1982.

[38] I. S. Gopal and C. K. Wong, "Minimizing the number of switchings in an SS/TDMA system," *IEEE Trans. Communications*, vol. 33, pp. 1497–1501, June 1985.

[39] N. M. Bhide, M. Mishra, and K. M. Sivalingam, "Scheduling algorithms for star-coupled WDM networks with tunable transmitter and tunable receiver architecture," *Photonic Network Communication*, pp. 219–234, Nov. 1999.

[40] E. Coffman, M.R. Garey, and D.S. Johnson, "An application of bin-packing to multiprocessor scheduling," *SIAM J. Computing*, vol. 7, pp. 1–17, Feb. 1978.