



Washington University in St. Louis

SCHOOL OF ENGINEERING & APPLIED SCIENCE

Design of Wavelength Converting Switches for Optical Burst Switching

Jeyashankher Ramamirtham and Jonathan Turner

WUCS-01-21

August 7, 2001

Department Of Computer Science,
Washington University in St. Louis,
Campus Box 1045,
One Brookings Drive,
St. Louis, MO 63130.

Design of Wavelength Converting Switches for Optical Burst Switching^{*}

Jeyashankher Ramamirtham, Jonathan Turner[†]

Abstract

Optical Burst Switching (OBS) is an experimental network technology that enables the construction of very high capacity routers, using optical data paths and electronic control. In this paper, we study two designs for *wavelength converting switches* that are suitable for use in optical burst switching systems and evaluate their performance. Both designs use tunable lasers to implement wavelength conversion. One is a strictly nonblocking design, that also requires optical crossbars. The second substitutes Wavelength Grating Routers (WGR) for the optical crossbars, reducing cost, but introducing some potential for blocking. We show how the routing problem for the WGR-based switches can be formulated as a combinatorial puzzle or game, in which the design of the game board corresponds to the pattern of interconnections used to join the input sections of the switch to the output sections. We use this to show how the interconnection pattern affects the performance of the switch, and to facilitate the design of interconnection patterns that yield the best performance. Our results show that for a typical switch configuration, the WGR-based design can deliver more than 87% of the throughput of a fully nonblocking switch.

1 Introduction

The transmission capacity of optical fibers has been increasing at a tremendous rate as a result of DWDM technology. Although terabit capacity IP routers based on electronics are now starting to appear, there remains a serious mismatch between the transmission capacity of DWDM fibers and the switching capacity of electronic routers. Since DWDM

^{*}This work is supported by the Advanced Research Projects Agency and Rome Laboratory (contract F30602-97-1-0273)

[†]J. Ramamirtham and J. Turner are with Applied Research Laboratory, Department of Computer Science, Washington University in St. Louis. E-mail: {jai,jst}@arl.wustl.edu

links are capable of supporting hundreds of channels operating at rates of 10 Gb/s each, it can take 5-10 equipment racks to hold the electronic line cards needed to terminate the channels from just a single fiber. Optical burst switching seeks to reduce the cost and complexity of these systems by replacing much of this electronics with optical components. OBS is a hybrid switching technology that uses electronics to control routing decisions, but keeps data in optical form as it passes through each OBS router. By exploiting the high channel counts of advanced WDM systems, it achieves excellent statistical multiplexing performance with little or no buffering. Further details on OBS can be found in [1].

In this paper we focus on the design of the basic switch elements that are used to construct large OBS routers. In particular, we study the design of the wavelength converting switches that are the key building block needed to implement these systems. Although, there have been a number of studies of optical packet switching in recent years [2, 3, 4], it is not yet clear how the required optical components can be implemented to make them cost-competitive with electronic alternatives.

Recent dramatic advances in tunable lasers have created new architectural options for wavelength converting switches and appear to hold considerable promise for the design of practical optical switching systems. This paper examines two architectures for wavelength converting switches for an optical burst switch. The first one uses Tunable Wavelength Converters along with optical crossbars and passive multiplexors and demultiplexors. The second design replaces the optical crossbars of the first design with passive Wavelength Grating Routers (also known as Arrayed Waveguide Grating Multiplexors or AWGMs). This option is very attractive because WGRs are relatively simple to fabricate, are inexpensive and consume no power. Unfortunately, the use of wavelength routers introduces the possibility of blocking. In this paper we evaluate the impact of this blocking on the statistical multiplexing performance of a switch element in an OBS router. Our results show that the WGR-based switch can achieve more than 87% of the throughput obtained with the fully nonblocking switch.

We show that the performance of a WGR-based switch is strongly dependent on the pattern of interconnections used to join the input and output sections of the switch. This is done by first formulating the routing problem in the WGR switch as a combinatorial puzzle or game, in which the design of the game board corresponds to the interconnection patterns. We use this correspondence to explore alternative designs and evaluate the performance of

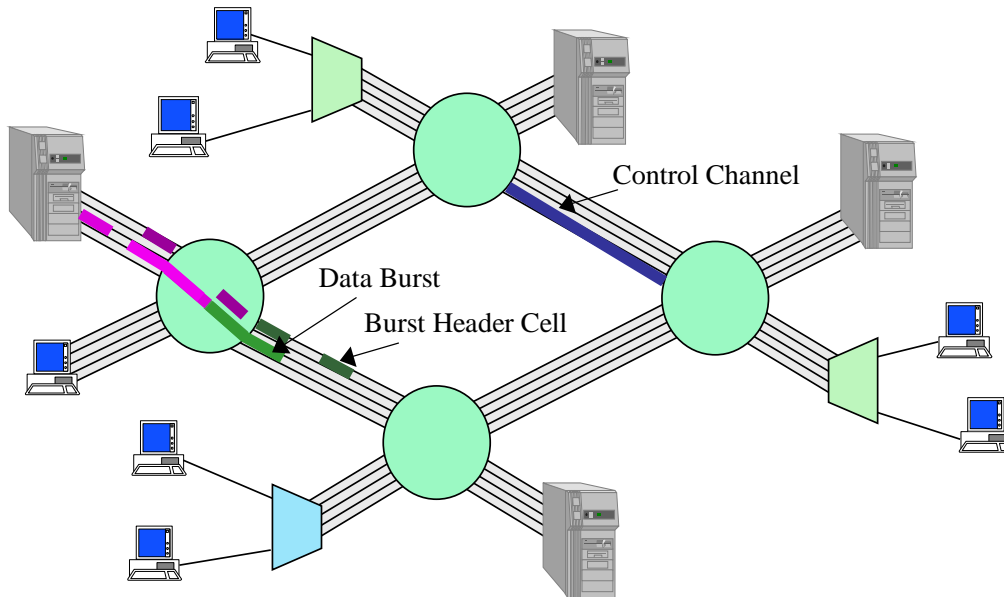


Figure 1: Burst Switching Concept

the resulting system, using simulation.

The rest of the paper is organized as follows. In Section 2, we give a brief overview of the optical burst switching concept and explain where this work fits in the OBS context. We present the two wavelength converting switch designs in Section 3. In Section 4, we show how to model the blocking in the WGR-based switch as a combinatorial puzzle or game. In Section 5, we present simulation results showing how the blocking characteristics of the switch designs affect the statistical multiplexing performance of a switch element in an OBS router. Finally, we present some concluding remarks in Section 6.

2 Burst Switching Architecture

The basic burst switching concept is illustrated in Fig. 1. The transmission links carry data on tens or hundreds of wavelength channels and user data bursts can be dynamically assigned to any of these channels by the OBS routers. One (or possibly several) channel on each link is reserved for control information that is used to control the dynamic assignment of the remaining channels to user data bursts. When an end system has a burst of data to send, an idle channel on the access link is selected and the data burst is sent on that channel. Shortly before the burst transmission begins, a *Burst Header Cell* (BHC) is sent on the control channel, specifying the channel on which the burst is being

transmitted and the destination of the burst. The OBS router, on receiving the BHC, assigns the incoming burst to an idle available channel at the outgoing link leading toward the desired destination and establishes a path between the specified channel on the access link and the channel selected to carry the burst. It also forwards the BHC on the control channel of the selected link, after modifying the cell to specify the channel on which the burst is being forwarded. This process is repeated at every router along the path to the destination. The BHC also includes an *Offset* field which contains the time between the transmission of the first bit of the BHC and the first bit of the burst, and a *Length* field specifying the time duration of the burst. The offset and length fields are used to time switching operations in the OBS routers, and the offset field is adjusted by the routers to reflect variations in the processing delays encountered in the routers' control subsystems. If a router does not have idle channels available at the output port, the burst can be stored in a buffer.

Reference [1] describes a scalable OBS router architecture consisting of a set of *Input/Output Modules* (IOM) that interface to external links and a multistage interconnection network of *Burst Switch Elements* (BSE). The interconnection network uses a Benes topology, which provides parallel paths between any input and output port. A three stage configuration comprising d port switch elements can support up to d^2 external links (each carrying many WDM channels). The topology can be extended to 5, 7 or more stages. In general, a $2k - 1$ stage configuration can support up to d^k ports. For example, a 5 stage network constructed from 8 port BSEs would support 512 ports. If each port carried 256 channels at 10 Gb/s each, the aggregate system capacity would be 1,310 Tb/s.

Input IOMs process the arriving BHCs, performing routing lookups and inserting the number of the output IOM into BHCs before passing them on. The BSEs use the output port number to switch the burst through to the proper output. Each of the components that does electronic processing on the cell keeps track of the time spent and updates the offset field in the BHC to maintain synchronization with the burst. Additional details can be found in [1].

3 WDM crossbar architectures

Each BSE in a burst switch requires a wavelength converting switch, capable of passing the switch from any input of the BSE's d input fibers to any of its d output fibers (Fig. 2). A BSE with $d = 8$ and $h = 256$ wavelengths would have an aggregate throughput of 2 Tbs, assuming 10 Gb/s per wavelength. Since bursts can arrive at unpredictable times,




Figure 2: Wavelength converting switch with d input/output fibers and h wavelength channels per fiber

a BSE must be able to switch bursts to different wavelengths on the output fibers, in order to provide acceptable statistical multiplexing performance at typical traffic intensities. Wavelength conversion technologies are discussed in [5, 6].

We consider two wavelength converting switch designs, both of which use optical modulators and tunable lasers to transfer a signal from an input wavelength to a tunable output wavelength. The first design also use optical crossbars to provide space division switching, while the second substitutes wavelength grating routers (WGR) for the crossbars.

3.1 Switch Based on Optical Crossbars

Fig. 3 shows the first wavelength converting switch design. Each of the d input sections has an optical demultiplexor that separates the different wavelength channels from each other before propagating through Tunable Wavelength Converters that quickly tune to any of the h output wavelengths. The wavelength converters are followed by $h \times d$ crossbars. Outputs of each crossbar are then connected to distinct passive multiplexors, which constitute the output section of the switch. The crossbars can be decomposed into $d \times d$ sections, followed by additional passive multiplexors, reducing the required size of individual crossbar components. Note that the crossbars implement a combination of switching and multiplexing, since there may be several signals on a given input fiber that are

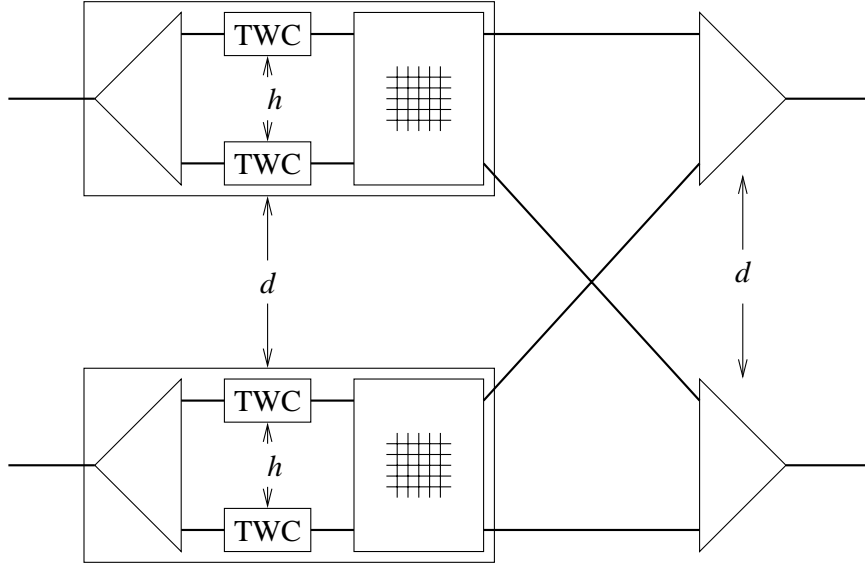


Figure 3: Wavelength converting switch using Tunable Wavelength Converters (TWC), Optical Crossbars and Passive Multiplexers and Demultiplexers

propagated to the same output fiber. So long as they are converted to distinct wavelengths, these signals may share the same crossbar output, so long as the optical crossbar technology is designed to support this. SOA-based crossbar technology is capable of implementing this combined switching and multiplexing function.

To route an incoming burst to an output, the input wavelength is converted to any available wavelength at the required output, and the crossbar is configured to propagate the signal to the required output. The switching and multiplexing capability of the crossbar ensures that there is no blocking, so long as there is an available wavelength on the selected output. Because burst arrival is unpredictable, there will be times in an OBS router when no output wavelength is available for an arriving burst. In routers with no internal buffering, such bursts are discarded. Fortunately, with high channel counts, the probability of burst discards is very low.

3.2 Switch Based on Wavelength Grating Routers

An alternative design for a wavelength converting switch is shown in Fig. 4. This design uses a passive wavelength grating router (WGR) in place of the optical crossbars used in the previous design. Thus, the tunable wavelength converters are the only active components. Since the wavelength routers have h inputs and h outputs, h/d fibers connect each input section with each output section. For $h = 256$ and $d = 8$, there will be 32 fibers connecting

each input section with each output section. In this design, the tunable wavelength converters serve two purposes. First, they are used to switch signals to distinct output wavelengths, to avoid wavelength conflicts on the output links. Second, in combination with the WGRs, they provide the required space switching. By tuning the laser to one wavelength in the appropriate set of h/d wavelengths for the desired output, we can “steer” the signal to the desired output port. The implication of this dual role of the wavelength converters is that choice of output wavelengths is constrained, leading to the possibility of blocking in this wavelength converting switch design. That is, there may be situations where all of the wavelengths that can be used to get to a desired output are in use at the output causing blocking to occur, even when there are free wavelengths available on the outgoing link. In the next section, we show how the problem of selecting wavelengths in such switches can be more easily understood by reformulating it as a combinatorial puzzle or game.

4 Design of WGR-based switches

In this section we study the routing problem in WGR-based switches and show how the blocking performance of these switches is affected by the interconnection pattern used to join the input sections of the switch to the output




Figure 4: Wavelength switch using Tunable Wavelength Converters (TWC) and Passive Wavelength Grating Routers (WGR)

	λ_0	λ_1	λ_2	λ_3
I_0	O_0	O_1	O_2	O_3
I_1	O_1	O_2	O_3	O_0
I_2	O_2	O_3	O_0	O_1
I_3	O_3	O_0	O_1	O_2

Figure 5: Routing matrix of a 4×4 WGR

sections. First however, we briefly review the routing characteristic of wavelength grating routers.

4.1 Wavelength Grating Router properties

An $h \times h$ WGR is a passive static wavelength-routing device that provides complete connectivity between its inputs and outputs, by passively routing h^2 optical connections on N wavelengths [7]. The use of a WGR has several advantages including easy fabrication, commercial availability and relatively low cost. A WGR has a fixed cyclical-permutation-based routing pattern between its input and output ports. A connection at input i using wavelength k gets routed to the same wavelength on output $(i - k) \text{ modulo } h, \forall i, k \in [0, h - 1]$. The routing pattern for a 4×4 WGR is shown in Fig. 5. A connection at input I_2 using wavelength λ_3 gets routed to output O_1 and a connection at input I_3 using wavelength λ_0 gets routed to output O_3 .

4.2 Routing Multiple Channels Simultaneously

In this section, we show how the problem of simultaneously routing a set of channels through a WGR-based switch can be formulated as a combinatorial puzzle. This formulation makes it easier to understand the intrinsic structure of the problem, yielding insights that are useful in design and analysis.

We start with the puzzle version of the problem, which corresponds to the problem of routing a set of connections through a switch, all at the same time. The puzzle is played on a game board made up of dh^2 squares arranged in h columns and dh rows. The board is divided into d square *blocks* of h rows each. Each square has one of d different colors, with each row containing h/d squares of each color and each column containing h squares of each color. To setup the puzzle we place colored *tokens* beside some or all of the rows. A setup can include at most h tokens of any

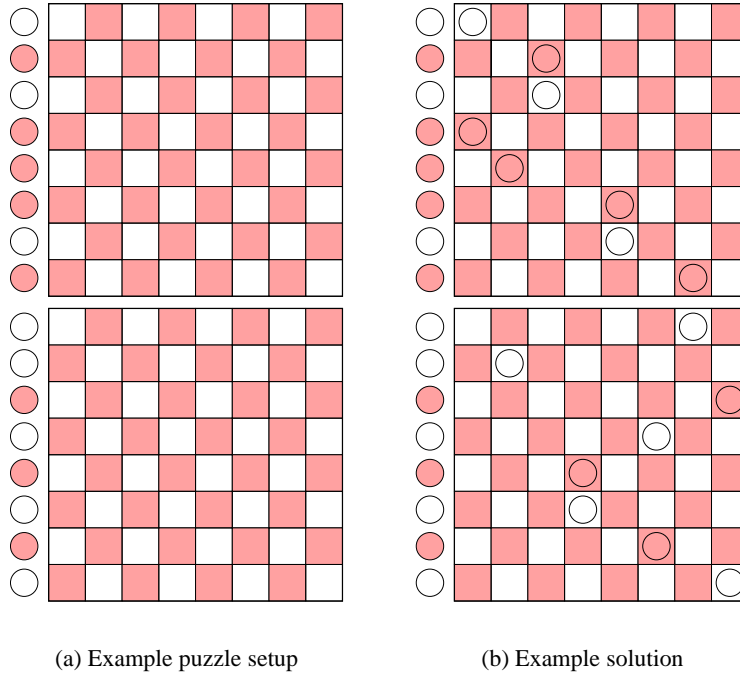


Figure 6: An example puzzle setup and solution

color. An example of a setup game board with $d = 2$ and $h = 8$ is shown in Fig. 6(a). To solve the puzzle, we must place each token on a square of the same color, in the row where the token was placed. The token placement must also satisfy the constraint that no two tokens of the same color be placed in the same column. An example solution to the puzzle is shown in Fig.6(b).

Each row in the puzzle corresponds to one of the h input channels on one of the d input fibers. More specifically, row i in block j of the game board, corresponds to input channel i of input fiber j . The color of the token that is placed by a row, corresponds to the output fiber that the corresponding input channel is to be switched to. More specifically, placing a token of color r on row i of block j corresponds to switching channel i of input fiber j to output fiber r . The columns of the array correspond to different output wavelengths. Placing a token in a particular column corresponds to choosing that output wavelength. The color of each square corresponds to the output that is reached if the wavelength converter for the input channel corresponding to that square's row is tuned to the wavelength corresponding to the column. So, placing a token of color r in column q of row i of block j corresponds to switching channel i of input fiber j to channel q of output fiber r . Note that the puzzle rule requiring that no two tokens of

the same color occupy the same column, corresponds to the requirement that no two input signals going to the same output fiber use the same wavelength.

In order to complete the correspondence between the puzzle and the routing problem, we note that within each block, the rows must have closely related color patterns, in order to model the routing characteristics of the WGRs. Specifically, the pattern of colors within each row can be obtained from the previous row's pattern by a cyclic rotation of one column. This relationship only holds within each block. There is no requirement that different blocks have similar color patterns. The color pattern for each block corresponds to the pattern of interconnections joining the input sections of the switch and the output sections. This is illustrated in Fig. 7 which shows two example configurations of a system with $d = 2$ and $h = 8$ and the corresponding game boards.

Whenever the puzzle has a solution, it means that there is a way to route the input signals to the output channels that are specified by the tokens placed by each row. If the puzzle does not have a solution, then there is no way to route all the channels simultaneously. If, for all possible puzzle setups, there is a solution, the switch is *rearrangeably non-blocking*. It is easy to see that the switch in Fig. 7a is not rearrangeably non-blocking, since the puzzle setup in which tokens of one color are placed in even-numbered rows and tokens of the other color are placed in odd-numbered rows, has no solution. On the other hand, this setup does have a solution when played on the game board in Fig. 7b.

4.3 Finding Good Game Boards

The design of a game board has a big influence on our ability to solve the puzzle. Since the game board design corresponds to the interconnection pattern of the switch, this means that the interconnection pattern affects the likelihood of blocking. The game board in Fig. 7a has many puzzle setups that have no solution, making it a poor design, from the perspective of the puzzle solver. What makes it a poor design is that many rows have exactly the same pattern of colors. This means that if tokens of the same color are placed in these rows, the number of columns they have to choose from is limited, and may be smaller than the number of tokens. This suggests that a good game board design will be one in which different rows have different patterns, and in particular, have as few columns as possible with squares of the same color.

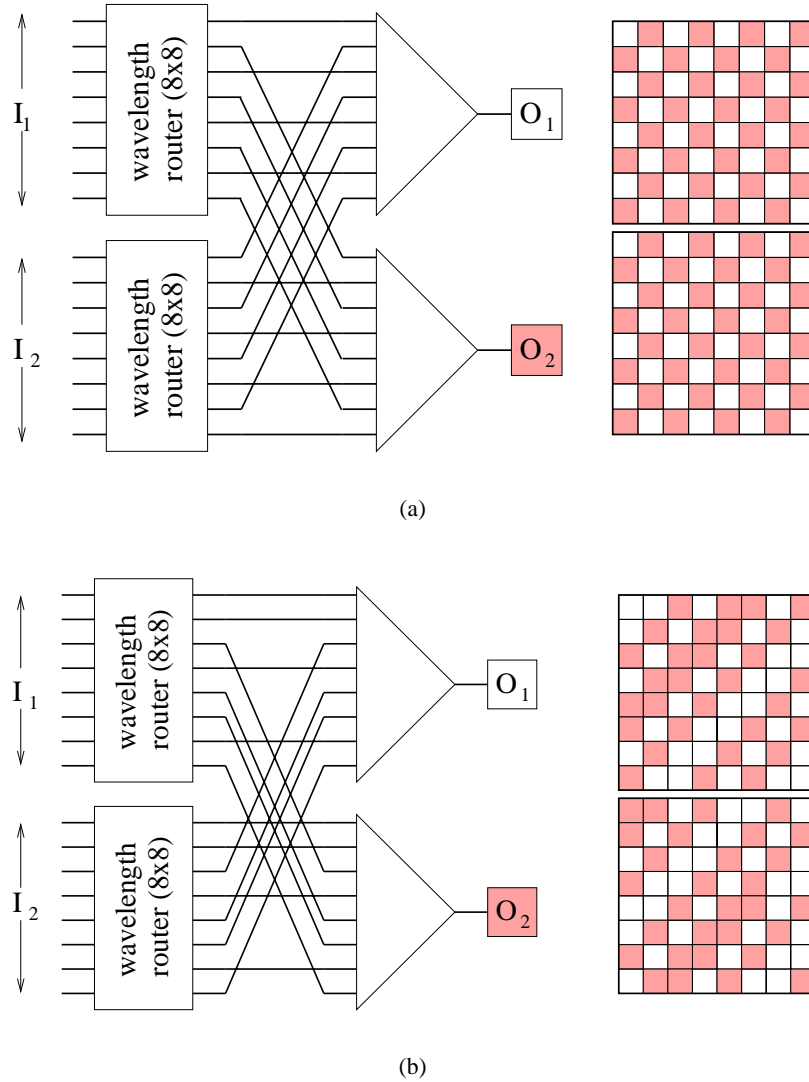


Figure 7: Two configurations and the corresponding game boards of a system with $d = 2$ and $h = 8$

For a given game board, define $C_j(i)$ be the set of columns with squares of color j in row i . We call this the j -cover of row i . Similarly, for any set of rows R , let $C_j(R) = |\cup_{i \in R} C_j(i)|$ be the j -cover of R . If for all colors j and all sets of rows R with $\leq h$ rows, the j -cover of R has at least $|R|$ columns, then the puzzle always has a solution, and hence the corresponding switch is rearrangeably nonblocking. This can be proved by considering tokens of each color separately and using the above property to show how a partial solution with $i < h$ tokens of a given color can be extended to a solution with $i + 1$ tokens of that color.

Unfortunately, there is no game board design that always has a solution. To see this, consider an arbitrary game

board and some color (call it “blue”). There are exactly h blue squares in any column of the game board, meaning that there are $dh - h$ squares that are not blue. If we select any h rows from among the $dh - h$ rows that do not have blue squares in the given column, then any puzzle setup that has blue tokens in these h rows is unsolvable, since none of the tokens can be placed in the selected column. Similarly, if we consider any $i \leq d - 1$ columns, there must be at least $(d - i)h$ rows that do not contain blue squares in any of these columns. So, any puzzle setup that has blue tokens in more than $h - i$ of these rows is unsolvable. These results make it clear that we cannot expect to construct a WGR-based switch that will guarantee our ability to place more than $h - d + 1$ tokens of the same color. Fortunately, the value of h is typically much larger than d for configurations of practical interest, which means that the degree of blocking implied by this limitation may be acceptable.

To assess the practical impact of blocking in WGR switches, we performed an experiment using random game board configurations. We selected random row sets of various sizes and determined the average size of the j -covers for all j . More precisely, for each size, we selected 10,000 row sets and computed the average size of the j -covers over all j and all sets of the given size. The results of this experiment are plotted in Fig 8. For the case of $d = 64$, we also show the size of the smallest j -covers found.

For the values of d of most practical interest (≤ 16), the j -covers are typically much larger than the number of rows until we get to the largest row sets. This suggests that blocking can be avoided if the number of signals going to a given output is limited to a value slightly smaller than h .

4.4 Routing Bursts in an OBS Router

The puzzle introduced above corresponds to the version of the routing problem in which we are asked to simultaneously route a whole set of connections. In burst switching systems however, bursts arrive at different times and must be routed through a switch without disturbing bursts that are already in progress. The problem of routing individual bursts can be formulated as a two player game, played on the same game board as the puzzle.

Lets call the first player the *blocker* and the second player, the *setter*. The blocker is given up to dh tokens of d different colors, with up to h tokens of each color. The blocker takes a turn by removing zero or more tokens from the board and placing a token beside some row of the board. The setter takes its turn by placing the token put down

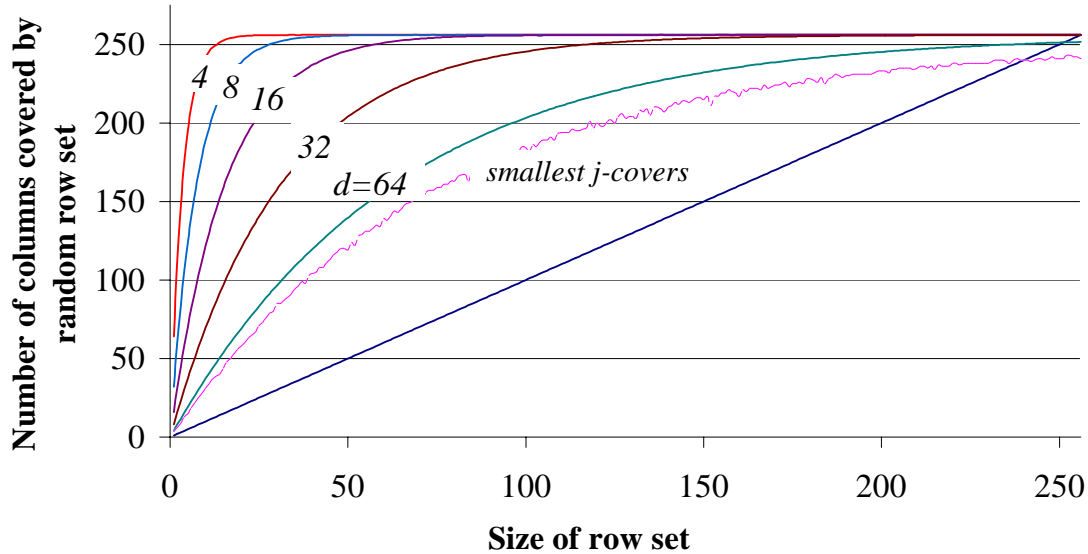


Figure 8: Number of columns covered by row sets of all sizes ($h = 256$)

by the blocker, in a square of the same color as the token in the selected row. When placing the token, the setter must not use any column that already contains a token of the same color. The blocker wins if the setter is not able to place the token on the board without violating the conditions. The blocker loses if the setter is able to keep the game going indefinitely. The switch is non-blocking if there is a winning strategy for the setter (that is a strategy that will keep the game going forever, regardless of how well the blocker plays).

Since a winning strategy for the setter would imply that the corresponding puzzle always has a solution, we cannot expect a winning strategy in versions of the game where the blocker has more than $h - d + 1$ tokens of any color. The existence of winning strategies in versions of the game where the number of tokens of each color is limited to at most $h - d + 1$ remains open.

In the next section, we study how the blocking characteristics of WGR switches affect the statistical multiplexing performance of Burst Switch Elements in OBS routers.

5 Simulation results

This section reports the results of simulations done to evaluate the performance of WGR-based switches in an OBS router. Here, we consider only the case of routers in which there are no buffers available to store bursts which can't




Figure 9: Error probabilities of different system configurations ($d = 8$)

be routed to the proper output without a wavelength conflict. Burst arrivals on each input channel were independent and each arriving burst was randomly assigned to a different output fiber. Burst lengths and the idle times between successive bursts on the same channel were exponentially distributed. The simulations used random interconnection patterns to join the input and output sections of the switch. Arriving bursts were assigned the first wavelength that would take them to proper output, that was not already in use at that output. In the game formulation, this corresponds to placing a token in the leftmost square of the right color, for which the column does not already contain a token of the same color.

The performance metric used is the fraction of arriving bursts that must be discarded. This is called the *burst rejection probability*. The burst rejection probabilities for systems with different values of d and h and varying loads are shown in Fig. 9. Also shown are the burst rejection probabilities for systems that use strictly nonblocking switches in place of WGR-based switches. For a system with $d = 8$ and $h = 256$, the rejection probability is 10^{-6} at a load of approximately 0.62 for the WGR-based switch and at a load of 0.75 for the strictly nonblocking switch. For systems designed to operate with a burst rejection ratio of 10^{-6} , the WGR-based switch can provide a throughput which is approximately 82% of what the nonblocking switch can provide.

The next set of results show how different game board configurations can affect system performance. Fig. 10 shows the result of using four different configurations for a system with $d = 8$ and $h = 256$. The first configuration corresponds to a “consecutive” interconnection pattern, wherein, each input block’s first h/d outputs are connected to the first output, the next h/d outputs are connected to the second output, and so on. The second configuration corresponds to a perfect shuffle between the wavelength router stage and the output side couplers. The third configuration corresponds to a randomly generated game board. The fourth configuration is a hand generated configuration where colors corresponding to any output in a row are distributed such that two rows have very little overlap and hence, the number of wavelengths available to reach a given output between a set of rows is increased. As can be seen from Fig. 10, the first two configurations perform poorly. This is because they do not try to maximize the number of wavelengths available to a subset of input rows uniformly. The fourth configuration performs slightly better than the random pattern and has a burst rejection probability of 10^{-6} at a load of 0.65. With this design, the WGR-based achieves 82% of the throughput that is achieved with the strictly non-blocking switch.

6 Conclusion

We have studied the performance of two wavelength converting switch designs for use in OBS routers. The first one uses optical crossbars and Tunable Wavelength Converters and is strictly non-blocking. The second design substitutes Wavelength Grating Routers for the optical crossbars, making it significantly less expensive. Although the blocking nature of the second design leads to higher burst rejection probabilities, the performance penalty is small enough to make it a viable alternative.

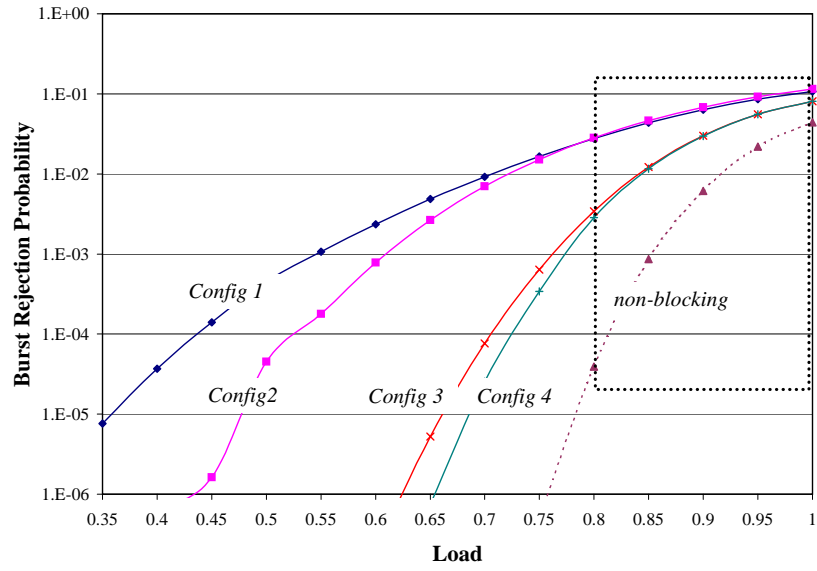
By formulating the routing problem as a combinatorial puzzle or game, we have been able to develop insights that facilitate the analysis and design of WGR-based switches. We have shown some basic limits to the nonblocking potential of WGR-based switches, but have also shown that by selecting the interconnection patterns appropriately, one can greatly improve their performance. Simulation results show that in practical switch system configurations, routers using WGR-based switches can achieve more than 87% of the throughput of routers using strictly non-blocking switches.

An interesting design option that needs to be explored is the use of buffering in OBS routers. In general,


buffering can be expected to improve the performance, and we expect it to have a larger impact on routers built using WGR-based switches, narrowing the performance gap further.

References

- [1] Jonathan S. Turner, "Terabit Burst Switching," *Journal of High Speed Networks*, 1999.
- [2] D. Blumenthal, P. Prucnal, J. Sauer, "Photonic packet switches: Architectures and experimental implementation," *Proc. IEEE*, vol. 82, pp. 1650-1667, Nov. 1994.
- [3] P. Gambini, *et al.*, "Transparent optical packet switching: Network architecture and demonstrators in the KEOPS project," *IEEE Journal on Selected Areas in Communication*, vol. 16, pp. 1245-1259, Sep. 1998.
- [4] Francesco Masetti, "System functionalities and architectures in photonic packet switching," *Photonic Networks*, Giancarlo Prati (editor), Springer Verlag 1997.
- [5] T. Durhuus, B. Mikkelsen, C. Joergensen, S.L. Danielsen, K.E. Stubkjaer, "All-optical wavelength conversion by semiconductor optical amplifiers," *Journal of Lightwave Technology*, vol. 14, pp. 942-954, June 1996.
- [6] S.J.B. Yoo, "Wavelength conversion technologies for WDM network applications," *Journal of Lightwave Technology*, vol. 14, pp. 955-956, June 1996.
- [7] Dragone, C., "An NxN optical multiplexor using a planar arrangement of two star couplers," *IEEE Photonic Technology Letters*, vol.6, pp. 812-815, Sept. 1991.



(a)



(b)

Figure 10: Error probabilities of different game boards ($d = 8, h = 256$)