Efficient Routing in Optical Networks*

Alok Aggarwal Amotz Bar-Noy† Don Coppersmith
Rajiv Ramaswami Baruch Schieber Madhu Sudan

IBM Research Division
T. J. Watson Research Center
Yorktown Heights, NY 10598

Abstract

This paper studies the problem of dedicating routes to connections in optical networks. In optical networks, the vast bandwidth available in an optical fiber is utilized by partitioning it into several channels, each at a different optical wavelength. A connection between two nodes is assigned a specific wavelength, with the constraint that no two connections sharing a link in the network can be assigned the same wavelength. This paper considers optical networks with and without switches, and different types of routing in these networks. It presents optimal or near-optimal constructions of optical networks in these cases and algorithms for routing connections, specifically permutation routing for the networks constructed here.

*This work was supported in part by a grant No. MDA 972-92-C-0075 from ARPA. An earlier version appeared at the ACM Symposium on Discrete Algorithms, 1994.

†Current address: Dept. of EE-Systems, Tel Aviv University, Israel. e-mail: amotz@eng.tau.ac.il.

1. Introduction

Fiber-optic networking technology using wavelength division multiplexing (WDM) offers the potential of building large wide-area networks capable of supporting thousands of nodes and providing capacities of the order of gigabits-per-second to each node in the network [Gre92, Ram93, CNW90]. In WDM optical networks, the vast bandwidth available in optical fiber is utilized by partitioning it into several *channels*, each at a different optical *wavelength*. Each wavelength can carry data modulated at bit rates of several gigabits per second.

In general, such a network consists of *wavelength routers* interconnected by point-to-point fiber-optic *links* (Figure 1). Each link can support a certain number of wavelengths. Wavelength routers are the analogs of switches in traditional networks. Each router is an endpoint of several optical links. Each link is connected to either an input port of the router or an output port. The router determines which incoming signal is to be transmitted to which outgoing link. There is a significant amount of variance in the functionality of a router — in terms of when it determines the routing pattern, how it determines the routing pattern and if it changes the signals it transmits in any fundamental way. The one common feature that all routers share is that they cannot route two signals on the same wavelength on the same outgoing link. We shall presently summarize the main categories of wavelength routers. In addition to routers and links, a network also consists of several end nodes. Each node is connected to a router of the network and consists of a tunable optical receiver and tunable optical transmitter. The transmitter can be tuned to transmit on any of the available wavelengths and the receiver can be tuned to receive on any of the available wavelengths.

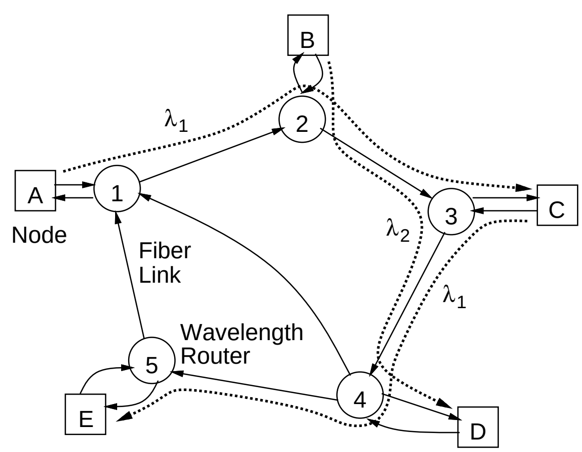

Figure 1: A WDM network consisting of wavelength routers interconnected by point-to-point fiber-optic links. Some of the wavelength routers have nodes attached to them that form the sources and destinations for network traffic.

Wavelength routers. The simplest form of a wavelength router is a *non-reconfigurable* router. In a non-reconfigurable router, the incoming to outgoing pattern is determined *statically* and can not be switched once the router is built. The main feature of such routers is the fact that the routing pattern can be a function of the wavelength of the incoming signal. Thus for each input port and each wavelength, the router associates a fixed set of output ports on to which it will transmit the incoming signal on the same wavelength (as the incoming signal). Figure 2 shows such a wavelength router. It is a passive (unpowered) device, and can be realized in many forms. The realization here uses passive wavelength demultiplexers and multiplexers.

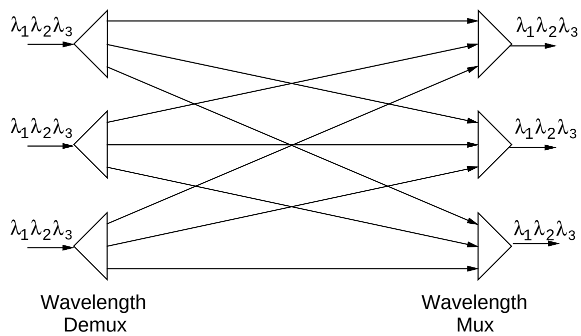

Figure 2: A non-reconfigurable (switchless) wavelength router with 3 ports capable of handling 3 wavelengths per port.

A second form of a wavelength router is the *wavelength-independent reconfigurable* switch. (In previous versions of the paper this switch was also called an *elementary* switch.) These switches are functionally identical to switches in the classical networks. Each switch can be dynamically reconfigured to change its input-output pattern. However, the input-output pattern is independent of the wavelength of the incoming signal. Thus for a fixed configuration of the switch, there is a fixed set of output ports on to which an incoming signal is transmitted independent of their wavelengths.

A third form of wavelength router which generalizes both the above forms is the *wavelength-selective reconfigurable* switch. (In previous versions of the paper this switch was also called a *generalized* switch.) As in the wavelength-independent case, the input-output pattern of the wavelength-selective switch can also be switched dynamically. However, in this case the routing pattern can also be a function of the wavelength of the incoming signal. Thus for a fixed configuration, for every input port and wavelength of the incoming signal, the switch associates a fixed set of output ports onto which it transmits this signal. It is clear that this switch can simulate either of the routers mentioned earlier. It is also true that a combination of non-reconfigurable routers and wavelength-independent reconfigurable switches can simulate a wavelength-selective reconfigurable switch. (See Figures 3(a) and (b) for an example of how a wavelength-selective switch is simulated by wavelength multiplexers and demultiplexers, and wavelength independent switches.)

Lastly, we mention that the literature has also considered a class of wavelength routers with an additional feature – that of wavelength conversion. Such routers are capable of changing the wavelength of an incoming signal before transmitting it to an outgoing optical. It is possible to consider a variant of all the above routers (reconfigurable and non-reconfigurable) with this

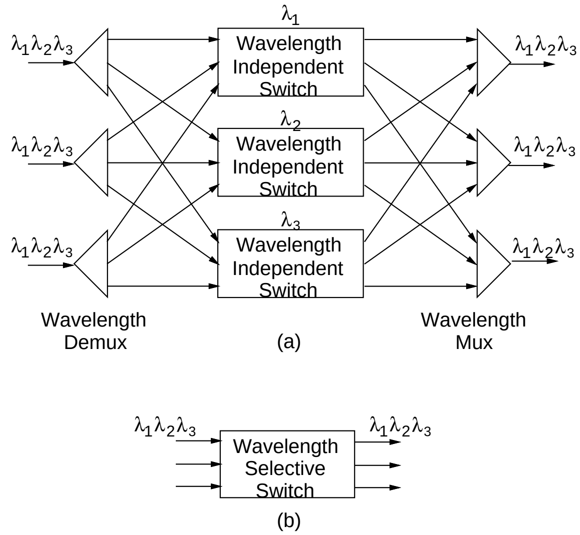

Figure 3: Two types of reconfigurable wavelength switches, each with 3 ports capable of handling 3 wavelengths per port. They can switch each wavelength at its input ports independent of the other wavelengths. The switches can be reconfigured to allow different interconnection patterns. (a) This router uses wavelength-independent switches in combination with passive wavelength demultiplexers and multiplexers. (b) This router uses a wavelength-selective switch.

additional feature. However in this paper we shall not be considering this type of wavelength router.

Classification of Networks. Based on the above classification of routers, we consider networks which have a subset of these routers available. The first class of networks that we consider are *non-reconfigurable*, or *switchless* networks. The only form of wavelength routers available in such networks are the non-reconfigurable routers. Switchless networks are of practical importance because the network components are passive and hence reliable, and moreover, all the control is done outside the network, making network control easier. Another motivation for studying such networks is that they form important components in the construction of *reconfigurable* networks which we describe next (Section 3).

The second class of networks, which we call *reconfigurable networks*, are networks which consist of some non-reconfigurable routers and some reconfigurable routers. Within the class of reconfigurable networks we consider two subcases: (i) networks that allow only wavelength-independent reconfigurable switches and non-reconfigurable routers, and (2) networks that allow wavelength-selective reconfigurable switches and non-reconfigurable ones.

In switchless networks, once the routing pattern is set, the only choice remaining is in selecting the wavelength at which each node transmits and the wavelength at which it receives. In reconfigurable networks, additional degrees of freedom are obtained by changing the settings of the switches. We assert that optical switches will be required to build large networks because the switchless network requires a large number of wavelengths to support even simple traffic patterns (as will be seen later in this paper).

Problems and parameters of interest. In an optical network several node pairs may request to be "connected". A *connection* between a pair of nodes is a path connecting the two nodes and a wavelength. A set of connections is *legal* if no two paths using the same wavelength overlap on a link (or an edge). (See, for example, Figure 1 which shows a connection from node A to node C on wavelength λ_1 and a connection from C to E also on λ_1. However the connection from B to D must be carried on a different wavelength, λ_2.) The primary task we are interested in is the construction of networks which allows for a fairly general class of connection requests to be legally connected. A second task is the task of deciding how to set up the connections satisfying a given collection of connection requests in the given network. For the networks we construct we end up solving this second task easily, though in general networks this problem may be much harder.

In general the number of wavelengths required is a key parameter that we seek to minimize. For non-reconfigurable networks, the number of routers and their degree are not as important and we shall ignore these parameters in this paper.

In the case of reconfigurable networks again the parameter of interest is the number of wavelengths, but this time the number of reconfigurable switches used and their degree also becomes important. In particular, it appears that the cost of constructing switches handling more than a fixed number of ports may be too high. Thus in this paper we restrict attention to bounded degree switches and analyze the number of switches used as a function of the number of available wavelengths.

1.1. Terminology

A *permutation network* is a network that can successfully route all sets of connections where each transmitter is connected to a single receiver and each receiver to a single transmitter.

A permutation network handles *connection establishment requests* and *connection termination requests*. A connection establishment request specifies the transmitter and receiver between which a connection is to be established. It is assumed that both the transmitter and the receiver are idle when a connection establishment request is initiated. In this case a permutation network must always be able to set up this connection. A connection termination request specifies a pair (transmitter, receiver) that are currently connected and terminates this connection. Following Beneš [Ben62], we distinguish between two types of *non-blocking permutation networks* (NBNs): *rearrangeably NBNs*, where existing connections *can* be rerouted to accommodate a new connection establishment request, and *wide-sense NBNs*, where existing connections *cannot* be rerouted to accommodate a new connection establishment request.

An *oblivious* routing scheme always uses the same wavelength to satisfy a given connection request, regardless of the other connections in the network. Oblivious schemes are clearly on-line schemes. (An on-line scheme is a scheme that does not require the prior knowledge of future requests.) In *partially oblivious* routing, the wavelength that can be used to satisfy a connection request must be chosen from a subset of the available wavelengths in the network.

The *congestion* of a routing algorithm is the maximum number of paths that go over a single edge in the network. The *dilation* of a routing algorithm is the maximum number of edges in a path used by the routing algorithm.

1.2. Previous Work

The simplest form of a switchless optical network is a broadcast star network, shown in Figure 4. In a star network, a transmission from a node is broadcast to all the nodes in the network. Clearly a star network with n nodes requires n wavelengths for permutation routing. Also it is sufficient to provide each node with a *fixed-tuned* transmitter at a wavelength different from the other nodes, and a *tunable* receiver in order to be able to route permutations. An alternative is to make the transmitters tunable and the receivers fixed-tuned.

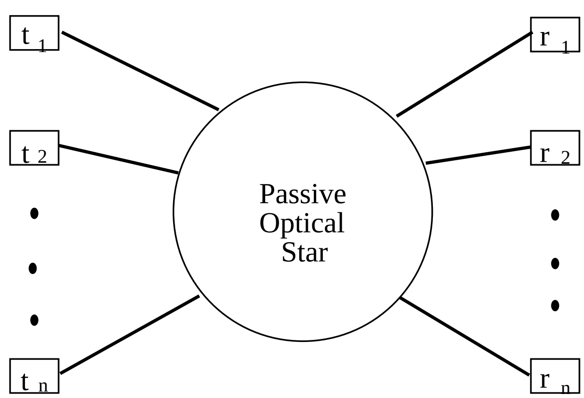

Figure 4: A broadcast star network.

However by using wavelength routers it is possible to route n connections using fewer than n wavelengths even in switchless networks. Barry and Humblet [BH92] showed that permutation routing in a switchless network with n nodes requires $\Omega(\sqrt{n})$ wavelengths. They also showed [BH92, BH93a] that oblivious permutation routing requires at least $\lfloor n/2 \rfloor + 1$ wavelengths and can be done using $\lceil n/2 \rceil + 2$ wavelengths.

In a reconfigurable network with w available wavelengths, Barry and Humblet [BH93a] showed that the number of wavelength-independent 2×2 switches required to support permutation routing is $\Omega(n \log(n/w^2))^\dagger$. For the special case in which the transmitters are fixed-tuned and the receivers are tunable, Pieris and Sasaki [PS93] showed that the number of wavelength-independent 2×2

†All logarithms in this paper are to the base 2.

switches required for permutation routing is $\Omega(n \log(n/w))$, and constructed such a network using $O(n \log(n/w))$ wavelength-independent switches.

Pankaj [Pan92] obtained bounds on the number of wavelengths required for permutation routing in certain network topologies using wavelength-selective switches. His network consisted of n wavelength-selective switches of fixed degree with each node being connected to a different router. For this model, Pankaj showed that $\Omega(\log n)$ wavelengths are required for permutation routing. He also showed that rearrangeably non-blocking permutation routing can be done with $O(\log^2 n)$ wavelengths and wide-sense non-blocking permutation routing with $O(\log^3 n)$ wavelengths in popular interconnection networks such as the shuffle exchange, de Bruijn, and hypercube networks.

1.3. Contributions of This Work

We present almost tight bounds for most of the problems considered in the earlier papers. Our results are summarized in Table 1.

Our first set of results are for switchless networks. We prove that oblivious permutation routing in switchless networks can be done using $\lceil n/2 \rceil + 2$ wavelengths, and prove that this is *optimal*; i.e., oblivious permutation routing in such networks requires $\lceil n/2 \rceil + 2$ wavelengths. The upper bound has been obtained independently by Barry and Humblet [BH92, BH93a]. They have also obtained a lower bound that is lower than ours by at least 1. We demonstrate the existence of a switchless permutation network using $O(\sqrt{n \log n})$ wavelengths. This result has been also obtained independently by Barry and Humblet [BH93b]. Both results give networks which are non-blocking in the wide-sense. We also give a polynomial time algorithm which dictates the tuning of the transmitters and receivers as the requests come online.

Unfortunately the above result is not a constructive one. We complement it with a constructive version which is only slightly weaker. Define $\xi(n) = 2^{(\log n)^{0.8+o(1)}}$. We show how to construct a switchless permutation network using $O\left(\sqrt{n}\xi(n)\right)$ wavelengths.[‡] We also provide non-trivial upper bounds for partially oblivious networks.

An important fact in our results for switchless networks is that the wavelength routers are not of fixed degree. In other words we ignore the complexity of switchless (passive) wavelength routers. In active switches however the complexity of a switch strongly depends on its degree, and all our constructions for reconfigurable networks use switches of fixed degree.

For reconfigurable networks with wavelength-independent switches, we show the existence of a wide-sense non-blocking network using $O(n \log \frac{n \log w}{w^2})$ wavelength-independent switches and construct a wide-sense non-blocking network using $O\left(n \log \frac{n\xi(w)}{w^2}\right)$ wavelength-independent switches. Clearly all of these results apply also for rearrangeably non-blocking networks.

For reconfigurable networks with wavelength-selective switches, we prove that any permutation network using w wavelengths requires $\Omega(\frac{n}{w} \log \frac{n}{w^2})$ wavelength-selective switches of constant degree.

[‡]Unfortunately, $\sqrt{n}\xi(n) < n$ only for $n > 10^{10}$.

Non-Reconfigurable (Switchless) Networks			
	NUMBER OF WAVELENGTHS		
	Lower Bound	Upper Bounds	
		Previous	Ours
Oblivious Routing	$\lceil \frac{n}{2} \rceil + 2$	$\lceil \frac{n}{2} \rceil + 2$ [BH92]	$\lceil \frac{n}{2} \rceil + 2$
Non-oblivious (existence)	$\Omega(\sqrt{n})$ [BH92]	$O(\sqrt{n \log n})$ [BH93b]	$O(\sqrt{n \log n})$
Non-oblivious (constructive)	$\Omega(\sqrt{n})$ [BH92]		$O(\sqrt{n}\xi(n))$
Partially oblivious ($k \geq 3$)	$\Omega(\sqrt{n})$ [BH92]		$O\left(n^{\frac{k+1}{2k-1}} k^{\frac{k}{2k-1}}\right)$
Reconfigurable Networks			
	NUMBER OF WAVELENGTHS		
$O(n)$ wavelength-selective switches	Lower Bound	Upper Bounds	
		Previous	Ours
Fixed topologies	$\Omega(\log n)$ [Pan92]	$O(\log^2 n)$ [Pan92] (rearrangeable) $O(\log^3 n)$ [Pan92] (wide-sense)	$O(\log n)$
Arbitrary topologies	$\Omega[c \min\{d, \sqrt{m}\}]$ (existential) (**)		$O[c \min\{d, \sqrt{m}\}]$
	NUMBER OF SWITCHES (w = NUMBER OF WAVELENGTHS)		
	Lower Bound	Upper Bounds	
		Previous	Ours
Wavelength-independent switches (existence)	$\Omega\left(n \log \frac{n}{w^2}\right)$ [BH92]		$O\left(n \log \frac{n \log w}{w^2}\right)$
Wavelength-independent switches (constructive)	$\Omega\left(n \log \frac{n}{w^2}\right)$ [BH92]		$O\left(n \log \frac{n\xi(w)}{w^2}\right)$
Wavelength-selective switches	$\Omega\left(\frac{n}{w} \log \frac{n}{w^2}\right)$		$O\left(\frac{n}{w} \log \frac{n}{w}\right)$

Table 1: Summary of results.

The results hold for both rearrangeable and wide-sense non-blocking cases unless specified otherwise. Notations used in the table: The number of nodes is n, c denotes congestion, d denotes dilation, m is the number of edges in the network, k denotes the number of wavelengths available for a connection, and $\xi(n) = 2^{(\log n)^{0.8+o(1)}}$.

(**) We show that there exist networks and connection requests that require $\Omega(c \min\{d, \sqrt{m}\})$ wavelengths to be connected.

We construct a permutation network with w wavelengths and $O(\frac{n}{w}\log\frac{n}{w})$ wavelength-selective 4×4 switches. Thus we can construct a permutation network with n wavelength-selective switches using only $O(\log n)$ wavelengths, an improvement over Pankaj's results [Pan92].

Next we consider networks with arbitrary topologies and arbitrary (not just permutation) connection establishment requests. We derive an upper bound on the number of wavelengths required for any routing scheme in terms of congestion and dilation for the given routing and the given network. We show that there exists a class of networks for which this bound cannot be improved.

The rest of the paper is organized as follows. Section 2 deals with non-reconfigurable (switchless) networks and Section 3 with reconfigurable networks (with switches). There are several open problems remaining to be solved; Section 4 gives a few such problems.

2. Non-Reconfigurable Optical Networks

In this section we consider non-reconfigurable (or switchless) optical networks. We use $\Lambda = \{\lambda_1,\ldots,\lambda_w\}$ to denote the set of available wavelengths. Observe that in such a network once we decide which wavelength a transmitter choose to transmit on, the set of receivers which can receive this signal is fixed. This allows us to model the network as a bipartite multigraph[§] $G(T,R,E)$ and a labelling function $\ell : E \to \Lambda$, where T is the set of *transmitters* and R is the set of *receivers* and for an edge e from a transmitter t to a receiver r, the label $\ell(e)$ denotes the wavelength which t can use to establish a connection to r. Since t may transmit to r using many possible wavelengths, there can be multiple edges between t and r. Thus two or more edges between a transmitter t and a receiver r will have different labels. A "tuning configuration" of the receivers and transmitters is formally described by a function $W : T \cup R \to \Lambda$. The tuning configuration is interpreted as follows. Every transmitter t transmits on the wavelength $W(t)$. Every receiver r receives on the wavelength $W(r)$. If e connects transmitter $t \in T$ to receiver $r \in R$, then whenever t transmits using wavelength $\ell(e)$, receiver r may receive this information only if it tunes to this wavelength. Moreover, since the network has no switches, *all* the receivers connected to t with edges labelled $\ell(e)$ receive t's message if they tune to this wavelength. Consequently, if a receiver r is tuned to a wavelength λ, then only *one* transmitter that is connected to r by an edge labelled λ may use this wavelength.

Note that once the graph G and the labeling are determined, the only choice remaining is in the tuning configuration of the transmitters and receivers. This motivates the following definition: A tuning configuration $W : T \cup R \to \Lambda$ is *valid* for a permutation $\Pi : T \to R$ in a network G, if for every $t \in T$, $W(\Pi(t)) = W(t)$ and for every pair of distinct transmitters $t_1, t_2 \in T$, either $W(\Pi(t_1)) \neq W(t_2)$ or if $W(\Pi(t_1)) = W(t_2) = \lambda$, then there is no edge labelled λ between $\Pi(t_1)$ and t_2 in G.

We consider the problem of constructing a non-blocking network using a minimum number of

[§]A multigraph is a graph with multiple edges allowed between nodes.

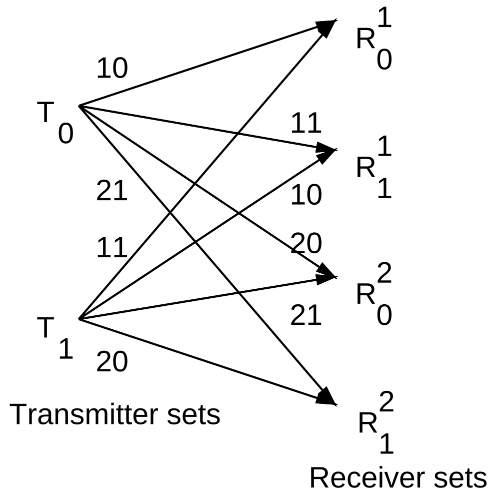

Figure 5: The construction of G for the case $b = 2, k = 2$. The edges are labelled with the wavelengths used for communication.

wavelengths.

We remark that the assumption that G is bipartite is made only to make the presentation clearer. We can achieve the same results for networks where some of the nodes are both transmitters and receivers. In this paper we consider the special case of $|T| = |R| = n$. However, most of the results can be extended to the case where $|T| \neq |R|$.

2.1. Rearrangeably Non-Blocking Networks

Barry and Humblet [BH92] proved that any non-reconfigurable rearrangeable NBN requires at least $(1 + \epsilon(n))\sqrt{n/e}$ wavelengths, where e is the base of the natural logarithm, and $\epsilon(n)$ goes to zero faster than $(\ln \sqrt{n})/\sqrt{n}$. We show that there exists a rearrangeably NBN that uses $O(\sqrt{n \log n})$ wavelengths. We show how to *construct* a network that uses $\sqrt{n}\xi(n)$ wavelengths (recall that $\xi(n) = 2^{(\log n)^{0.8+o(1)}}$).

In our upper bounds, the network has the following structure. The transmitter set T is partitioned into b blocks $T_0, \ldots, T_{b-1}$ each of cardinality either $\lceil n/b \rceil$ or $\lfloor n/b \rfloor$, where b is a parameter to be fixed later. (For clarity of exposition we from now on omit the $\lceil . \rceil$ $\lfloor . \rfloor$ operators.) The receiver set R is partitioned k times, with k to be fixed later. Each partition $1 \leq i \leq k$ partitions R into b blocks $R^i_0, \ldots, R^i_{b-1}$. (The cardinality of each such block may vary.) Our construction will use $w = b \cdot k$ wavelengths, denoted $\lambda_{i,j}$, for $1 \leq i \leq k$ and $0 \leq j \leq b-1$. The edges of the network are labelled as follows: for $1 \leq i \leq k$, $0 \leq j \leq b-1$, and $0 \leq a \leq b-1$, all the transmitters in T_a are connected to all the receivers in $R^i_{(a+j) \bmod b}$ by edges labelled $\lambda_{i,j}$. The construction for the case $b = 2, k = 2$ is illustrated in Figure 5.

The construction above has the following two properties:

G1 The transmitters are partitioned into b blocks where transmitters in each block are identically connected to all of the receivers.

G2 For any wavelength λ, if transmitters t_1 and t_2 belong to different blocks, then the set of receivers connected to t_1 by edges labelled λ is disjoint from the set of receivers connected to t_2 by edges labelled λ.

To get a rearrangeably NBN, it is necessary and sufficient to construct a network such that for any permutation $\Pi = \pi(1), \ldots, \pi(n)$, there is a way to tune the transmitters and receivers such that the connection requests $(t, \pi(t))$, for $1 \leq t \leq n$, are satisfied. A given tuning satisfies these connection requests if the following two properties are satisfied for all $1 \leq t \leq n$: (i) Both t and $\pi(t)$ are tuned to the same wavelength λ, and there exists an edge e connecting t to $\pi(t)$ with $\ell(e) = \lambda$. (ii) For all transmitters $t' \neq t$ such that there exists an edge e' labelled λ connecting t' to $\pi(t)$, t' is not tuned to λ.

Consider a permutation $\Pi = \pi(1), \ldots, \pi(n)$ that is to be routed. Property [G2] of our network implies that we can tune the transmitters of each block independently from the transmitters of other blocks. This is because transmitters from different blocks do not interfere. Property [G1] implies that in order to route Π, for any block of transmitters T_i, we have to use n/b different wavelengths. For this, the n/b destination receivers of the transmitters in T_i have to belong to n/b different blocks. Note that these blocks may belong to different partitions.

Given the network G, define the bipartite graph $H(S, B, F)$, where S, the *input* set, corresponds to the set of receivers (i.e., $S = R$), and B, the *output* set, corresponds to the set of blocks of receivers (i.e., $B = \{R_j^i : 1 \leq i \leq k, 0 \leq j \leq b-1\}$). A node $r \in S$ is connected by an edge in F to $R_j^i \in B$ if and only if the corresponding receiver r belongs to the corresponding block R_j^i. This construction is illustrated in Figure 6.

Observe that the graph H is characterized by the following two properties.

H1 The degree of each node in S is at most k.

H2 For $1 \leq i \leq k$, each node in S is connected to exactly one node from the set $\{R_0^i, \ldots, R_{b-1}^i\}$.

Theorem 1: *A network $G(T, R, E)$ with Properties [G1] and [G2] is non-blocking iff the corresponding graph $H(S, B, F)$ has the following* matching *property:*

H3 *The subgraph induced by any subset of n/b receivers and their neighbors in B contains a matching of cardinality n/b.*

Proof: We first prove that if $G(T, R, E)$ is non-blocking then $H(S, B, F)$ has the matching property. Let $S' \subset S$ be a set of cardinality at most n/b. Then consider a permutation Π from T to R

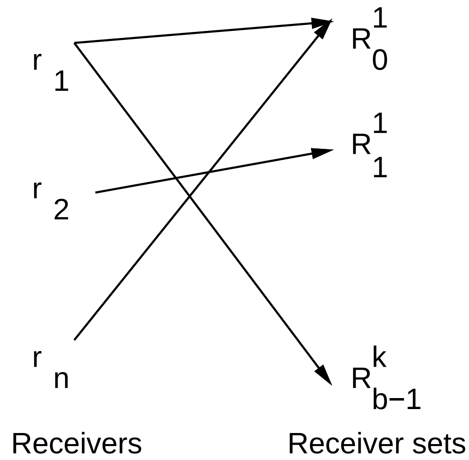

Figure 6: The construction of H.

which maps only transmitters in the set T_0 to elements of S'. (Such a permutation exists since $|T_0| = n/b \geq |S'|$.) Since G is non-blocking there exists a tuning configuration of the transmitters and receivers which is valid for Π. Now consider the assignment to the vertices of S' which assigns to a vertex $r \in S'$ the vertex R^i_j if the receiver $r \in R$ is tuned to the wavelength $\lambda_{i,j}$. Based on the validity of the tuning we have that no two vertices in S' are assigned the same vertex R^i_j in B. Furthermore, since the tuning assigns a wavelength $\lambda_{i,j}$ to a receiver r only if $r \in R^i_j$, we have that r is adjacent to its assigned vertex in the graph H. Thus the assignment above constitutes a matching in H.

We now prove the other direction, i.e., $G(T, R, E)$ is non-blocking if $H(S, B, F)$ has the matching property. Given a permutation $\Pi : T \to R$, we wish to construct a tuning configuration W that is valid for Π. Our strategy will be to construct a tuning of the transmitters such that if we extend it to the receivers by the mapping $W(\Pi(t)) = W(t)$, then the tuning is valid. Recall that Property [G2] of the construction implies that transmitters of each block can be tuned independently from the transmitters of other blocks. Fix a block of transmitters T_a. The destination receivers of the transmitters in T_a may be any subset of n/b receivers. Thus for any subset of n/b receivers, the receivers must belong to different blocks. By the definition of H, this translates to the matching property [H3]. $\square$

In the rest of this section we prove the existence of a graph H with properties [H1], [H2], and [H3] for $b = \sqrt{n/\log n}$ and $k = 4\log n$. Then, we show how to construct such a graph with $b = 2\sqrt{2n}$ and $k = 2\xi(n)$.

The results of [FFP88] imply the existence of a graph H having all three properties in which $b = \sqrt{n}$ and $k = O(\log n)$. The following theorem improves this result by a factor of $\sqrt{\log n}$.

Theorem 2: *There exists a graph $H(S, B, F)$ with properties [H1], [H2], and [H3] in which $b = \sqrt{n/\log n}$ and $k = 4\log n$.*

Proof: We construct $H(S, B, F)$ probabilistically as follows. Let $|S| = n$, and let B be partitioned into k blocks $B_1, \ldots, B_k$ of cardinality b each. We let each vertex in S pick k neighbors – one in each B_i independently and at random. We now analyze the probability that this graph has the matching property: i.e., any subset of up to n/b vertices from S is contained in some matching.

By Hall's Theorem [Hal35], such a matching does not exist if and only if there exists a set A of at most n/b vertices from S such that $|N(A)| < |A|$, where $N(A)$ denotes the set of neighbors of A. Since the degree of every vertex is k, it suffices to consider sets A of cardinality it at least k. Let α and $a_1, \ldots, a_k$ satisfy the conditions: $k \leq \alpha \leq n/b$ and $\sum_i a_i < \alpha$. Fix a set $A \subset S$ and sets $A_i \subseteq B_i$ such that $\alpha = |A|$ and $a_i = |A_i|$. The probability that $N(A) \subseteq \cup_i A_i$ is at most $\prod_{i=1}^{k} (\frac{a_i}{b})^\alpha$. Thus, the probability that there exist A and A_i's of cardinality α and a_i respectively such that $|N(A)| < |A|$ is at most

$$\binom{n}{\alpha} \prod_{i=1}^{k} \binom{b}{a_i} \left(\frac{a_i}{b}\right)^\alpha \leq \left(\frac{ne^2\alpha^{k-2}}{k^{k-1}b^{k-1}}\right)^\alpha .$$

Thus if $kb \geq cn/b \geq c\alpha$ for some constant $c \geq 2$ and $k = 4\log_c n$, then this probability goes to zero as $n^{-\Theta(\alpha)}$. For a fixed α, there are at most $\alpha^k = n^{4\log_c \alpha}$ choices for the a_i's. Since there are at most n/b choices for α, it follows that the probability that there exist α and a_i's such that this happens is bounded by $(n/b) \cdot \max_{k \leq \alpha \leq n/b}\{n^{-\Theta(\alpha)+4\log_c \alpha}\} = o(1)$.

Thus under the conditions $k = 4\log n$ and $b = \sqrt{n/\log n}$, we get that with a positive probability, $H(S, B, F)$ has the required three properties. □

Theorem 3: *A graph $H(S, B, F)$ with properties [H1], [H2], and [H3] in which $b = 2\sqrt{2n}$ and $k = 2\xi(2n)$ can be constructed.*

Proof: First, we define a concentrator.

Definition: An (x, y, ℓ)-*concentrator* with expansion α is a network with x inputs and y outputs such that every set of $t \leq \ell$ inputs expands to at least αt outputs. The *size* of the network is the number of edges and the *depth* of the network is the length of the longest path from an input to an output.

We use the following result from Wigderson and Zuckerman [WZ93].

Theorem 4 (WZ93): *For all x, there are explicitly constructible $(x, 2\alpha\sqrt{x}, \sqrt{x})$-concentrators with expansion α, depth 1 and size $\alpha x \cdot \xi(x)$.*

We now show how to apply Theorem 4. For our case, we set $x = 2n$ and $\alpha = 1$ and get that there exists a bipartite graph $H'(S', B', F')$, where $|S'| = 2n$, $B' = 2\sqrt{2n}$, and $|F'| = 2n \cdot \xi(2n)$ with the desired matching property. However, graph H' does not satisfy Properties [H1] and [H2]. We modify H' so that it satisfies these two properties. First, we consider a subgraph of H' which excludes all input nodes whose degree is more than two times the average degree in H'. Specifically, the degree of each input node in this subgraph is at most $2\xi(2n)$. Clearly, this subgraph still has the desired matching property. Because the size of the original graph H' is $2n \cdot \xi(2n)$, there are at least n input nodes in this subgraph. Next, we duplicate each output node $2\xi(2n)$ times and split the neighbors of each output node among the copies as follows. We number the edges outgoing from each input node with the numbers 1 to $2\xi(2n)$. Now, the first copy will have as edges the subset of the edges of the original node numbered 1, the second will have the subset numbered 2, and so forth. It is easy to see that the resulting graph has all the three properties. □

2.2. Wide-Sense Non-Blocking Networks

In this subsection we apply the above results to wide-sense NBNs.

Definition: A connection request is *one-sided* if it specifies only an input. It is satisfied by connecting the input to *any* of its neighboring outputs. A network H is wide-sense one-sided NBN if it can satisfy any sequence of *one-sided* connection and termination requests without rerouting.

Definition: A network H is a-limited wide-sense one-sided NBN if it can satisfy any sequence of requests in which at most a transmitters are connected simultaneously.

We show that to get a wide-sense NBN G, it is sufficient to make the corresponding graph H n/b-limited wide-sense one-sided NBN.

Theorem 5: *A network $G(T, R, E)$ with Properties [G1] and [G2] is wide-sense NBN if and only if the corresponding graph $H(S, B, F)$ is n/b-limited wide-sense one-sided NBN.*

Proof: The proof is similar to the proof of Theorem 1. We first prove that if $G(T, R, E)$ is wide-sense NBN then $H(S, B, F)$ is n/b-limited wide-sense one-sided NBN. We have to show how to satisfy in H any sequence of one-sided connection and termination requests as long as no more than n/b connections are active simultaneously. We do this by converting a sequence of one-sided connection and termination requests in H to a sequence of connection and termination requests in G, and then identify the matches in H with the wavelengths used to satisfy the requests in G. Consider a request to connect vertex $s \in S$ in H. Recall that s corresponds to a receiver in G. We associate with it a request to connect one of the transmitters in T_0 to the receiver s. (Since no more than n/b connections are active, and because of the way we convert termination requests, we are guaranteed to find a transmitter in T_0 that is not active.) Suppose that the wavelength used to satisfy this connection request in G is $\lambda_{i,j}$. Then, we match s to R^i_j. As in the proof of Theorem 1 it can be argued that since G is an NBN, R^i_j must be unmatched in H. A request to terminate the connection of vertex s in H, translates to a termination request of the connection that requested receiver s in G.

We now prove the reverse direction, i.e., If $H(S, B, F)$ is n/b-limited wide-sense one-sided NBN then $G(T, R, E)$ is wide-sense NBN. Consider a sequence of connection requests in G. Recall that Property [G2] of our construction implies that transmitters of each block can be tuned independently from the transmitters of other blocks. Thus is suffices to consider only the connection requests in which all transmitters are from some fixed block T_a. The destination receivers of the transmitters in T_a at any given time, may be any subset of at most n/b receivers. Thus, in order to satisfy any sequence of requests for this block in G, at any given time, all the receivers connected to transmitters in T_a must belong to different blocks. By the definition of H, this translates to the property that H must be n/b-limited wide-sense one-sided NBN. $\Box$

The following theorem is from [FFP88].

Theorem 6 (FFP88): *A network $H(S, B, F)$ is a-limited wide-sense one-sided NBN if every set X of inputs of cardinality at most $2a$ has at least $2|X|$ neighbors.*

We remark that by following the proof in [FFP88] we can actually prove that for our special case we may weaken the property, and consider only sets of cardinality at most a.

We conclude that the graph $H(S, B, F)$ is n/b-limited wide-sense one-sided NBN if it has the following property.

H4 Every set X of inputs of cardinality at most $2n/b$ has at least $2|X|$ neighbors.

Note that Property [H4] is stronger than Property [H3].

Theorem 7: *There exists a graph $H(S, B, F)$ with properties [H1], [H2], and [H4] in which $b = \sqrt{n/\log n}$ and $k = 10 \log n$.*

Proof: The proof is the same as the proof of Theorem 2. The only difference is in the value of the constants. $\Box$

Theorem 8: *A graph $H(S, B, F)$ with properties [H1], [H2], and [H4] in which $b = 4\sqrt{n}$ and $k = 2\xi(4n)$ can be constructed.*

Proof: We follow the construction given in the proof of Theorem 3. We use an explicit construction of $(4n, 8\sqrt{n}, 2\sqrt{n})$-concentrators with expansion 2, depth 1, and size $8n \cdot \xi(4n)$, and extract from it a graph with properties [H1], [H2], and [H4] in which $b = 8\sqrt{n}$ and $k = 2\xi(4n)$. $\Box$

There is one problem with our construction. Any algorithm that decides how to tune the transmitters and receivers appears not to be polynomial. Borrowing terminology from [FFP88] we have to maintain the maximum *critical* set of inputs. For this, it appears that after each request, we have to check all subsets of idle inputs. Below, we show how to alleviate this problem. We remark that it is not always easy to get polynomial decision algorithms for the routing questions based on the existential results. Arora, Leighton and Maggs [ALM90] give one such algorithm for their routing problem, other instances may be found from their paper. Our technique seems to be different from the previous methods.

We get a polynomial decision algorithm by strengthening the properties H has to satisfy.

Theorem 9: *A network $H(S, B, F)$ is a-limited wide-sense one-sided NBN and has a polynomial time decision algorithm if for every set X of inputs of cardinality at most a, even after we arbitrarily erase half of the edges adjacent to each input in X, X has at least $2|X|$ neighbors.*

Proof: At any stage, let A denote the subset of S corresponding to matching requests. For $v \in A$, let $T^*(v)$ denote its matched vertex in B and let $T^*(A)$ denote the set of all matched vertices in

B. The algorithm maintains: (i) a *critical set* C which contains the set A, and (ii) a match $T(v)$ for every vertex $v \in C$, such that T is an extension of T^*. The sets C, and $T(C) = \{T(v) | v \in C\}$ satisfy the following invariants.

Invariant 1: For each $x \in C \setminus A$, $|N(x) \cap T(C)| \geq \frac{1}{2}|N(x)|$; i.e., for each vertex x in C that is not in A, at least half of the neighbors of x are in $T(C)$.

Invariant 2: For every $x \notin C$, $|N(x) \cap T(C)| \leq \frac{1}{2}|N(x)|$; i.e., at least half of the neighbors of each of the vertices not in C is outside $T(C)$.

Invariant 3: $|T(C)| = |C| \leq 2|A| - 1$.

We remark that the critical set above is not equivalent to the critical set in [FFP88], though it attempts to capture the same set. We also remark that critical set is not unique and is a function of T.

We now show how the algorithm satisfies connection and termination requests, maintaining these invariants (under the assumption that these invariants held prior to these requests).

Case 1: A connection request for $x \in C$. In this case the algorithm matches x to $T(x)$. The set C and the map T remain unchanged. It is easy to see that invariants (1)–(3) above still hold.

Case 2: A connection request for $x \notin C$. In this case the algorithm picks a tentative match to x, denoted $t(x)$, outside $T(C)$. Such a neighbor must exist due to Invariant (2) above. Now, we compute the new inputs that have to be added to C. This is done incrementally. Let D be the current set of new inputs (D is initialized to $\{x\}$), and let D' be the set of outputs tentatively matched to these inputs (D' is initially $\{t(x)\}$). While there exists an y outside $D \cup C$ such that more than half of its neighbors are in $D' \cup T(C)$, find an output y' outside $D' \cup T(C)$ such that $D \cup \{y\}$ can be matched to $D' \cup \{y'\}$, and add y to D and y' to D'.

This is done as follows. First, find a matching M_1 of $D \cup \{y\}$ to outputs outside $T(C)$. Observe that as long as $|D| < a$, it follows from Invariant (2) and the property of H asserted in the theorem that $D \cup \{y\}$ satisfies the conditions of Hall's Theorem [Hal35], and hence such a matching exists. (In the claim below we prove that if such a vertex y exists, then $|C \cup D| < 2|A| + 1$, or $|D| \leq |A|$. Since we deal with a connection request $|A| < a$, and thus also $|D| < a$.) Consider the graph given by the union of M_1 and the matching M_0 of D to D'. The connected component of this graph that includes y must be an odd path that starts with y and ends with an output y' not in D'. The links of the path alternate between M_0 and M_1. The rest of the components must be either even paths or even cycles with alternating links. We construct the desired matching as follows: all the vertices from D that are in the "even" components are matched using the edges from M_0; vertex y, and all other the vertices from D that are in the "odd" component are matched using the edges from M_1. It is easy to see that the new matching matches $D \cup \{y\}$ to $D' \cup \{y'\}$. We update D to be $D \cup \{y\}$, D' to be $D' \cup \{y'\}$, and $t(x)$ to be the mate of x in the new matching.

We repeat this step till no such vertex y is found. Then, we update our matching $T^*(x)$ to be the final $t(x)$, the new set A to be $A \cup \{x\}$, the new set C to be $C \cup D$, and the new set C' to be $C' \cup D'$. It is clear that the invariants (1) and (2) hold upon termination. It remains to show that the procedure does eventually terminate and that when it does Invariant (3) above holds. The following claim shows that termination must occur when $|C \cup D| \leq 2|A|+1$. In addition of proving termination, since the size of A is incremented by one upon termination, this shows that Invariant (3) holds after the sets are updated.

Claim: If $|C \cup D| = 2|A|+1$ then there is no input vertex y outside $C \cup D$ such that more than half of the neighbours of y are in $T(C) \cup D'$.

Proof: To obtain a contradiction suppose that such a vertex y exists. We get that more than half of the neighbors of every input in $I = (C \setminus A) \cup (D \setminus \{x\}) \cup \{y\}$ are in a subset of outputs of size $2|A|+1$. Since $|A| \leq a-1$, $|I| = |C| - |A| + |D| - 1 + 1 = 2|A| + 1 - |A| = |A| + 1 \leq a$. However, the cardinality of the neighbors set of I is $|T(C) \cup D'| = |C \cup D| = 2|A|+1 < 2|I|$; a contradiction to the property of H stated in the theorem. $\square$

CASE 3: A termination request for x. Let x be an input for which a connection is terminated and let $T(x)$ be its matched output. We now need to show how to update C so that invariants (1)–(3) are satisfied. The obvious solution would be to retain the old C, but this may violate Invariant (3). Instead, we construct the set C "from scratch", as follows. Initially, we set C to be A, and $T(C)$ to be $T^*(A)$. Incrementally, we grow C similar to the previous case: As long as there exists an input $y \notin C$ such that more than half the neighbours of y are in $T(C)$, then add y to C. Note that y must have been in C also before the termination and thus $T(y)$ is defined. We add $T(y)$ to $T(C)$. This procedure is repeated until no such vertex y is found. Again, it is clear that if the process terminates invariants (1) and (2) hold. An argument similar to that in the claim above, shows that the process indeed terminates and that Invariant (3) holds as well. $\square$

We conclude that the graph $H(S, B, F)$ is n/b-limited wide-sense one-sided NBN and has a polynomial time decision algorithm if it has the following property.

H5 For every subset X of inputs of cardinality at most n/b, even after we arbitrarily erase half of the edges adjacent to each input in X, subset X has at least $2|X|$ neighbors.

Note that this property does not seem to imply Property [H4].

Theorem 10: *There exists a graph $H(S, B, F)$ with properties [H1], [H2], [H3], and [H5] in which $b = \sqrt{n/\log n}$ and $k = 16 \log n$.*

Proof: We follow the construction given in the proof of Theorem 2. The construction would yield a graph that does not have Property [H5] if there exists a set A of a vertices from S of cardinality at most n/b, and a set T of vertices from B, with cardinality less than $2a$, such that for each input w in A half the neighbours of w are in T. For a fixed size a, we upper bound the probability that sets $A \subset S$ of size a, and $T \subset B$ of size at most $2a$ with the above property exist by:

$$\binom{n}{a} \binom{k}{k/2}^a \cdot \prod_{i=1}^{k/2} \binom{b}{4a/k} \left(\frac{4a/k}{b}\right)^a \leq \left(\frac{ne^3 2^{2k-4} a^{k/2-3}}{k^{k/2-2} b^{k/2-2}}\right)^a .$$

Thus if $kb \geq cn/b \geq ca$ for some constant $c \geq 32$ and $k \geq 2\log_c n$, then this probability goes to zero as $n^{-\Theta(a)}$, and the probability that there exist a such that this happens can now be bounded by $o(1)$.

Thus under the conditions $k = 32\log n$ and $b = \sqrt{n/\log n}$, we get that with a positive probability, $H(S, B, F)$ has the required four properties. $\square$

We remark that a similar construction is used in [BRSU93] to obtain efficient routing in "classical" networks.

2.3. Oblivious and Partially Oblivious Routing

In this section we assume that whenever a transmitter has to communicate with a receiver it must use one out of a fixed number k of wavelengths. In previous sections we considered the case $k = w$, where w denotes the total number of wavelengths available. Here, we consider the case $k < w$. The case $k = 1$ is called the *oblivious* routing problem since there is no freedom in choosing wavelengths. Note that this implies that G is a graph rather than multigraph. The case $k \geq 1$ is called the *partially oblivious* routing problem. In this case, G is a multigraph with bounded multiplicity.

An oblivious routing network can be described by an $n \times n$ matrix M. The entry $M(i,j)$ in the matrix is an integer in the range $1, \ldots, w$ where w is the total number of wavelengths in the solution. The entry $M(i,j)$ indicates that i transmits to j using wavelength $M(i,j)$ in any permutation $\pi(.)$ for which $\pi(i) = j$.

Lemma 11: *Let the matrix M be a solution to the oblivious routing. If $\lambda = M(i,j) = M(i',j')$ for $i \neq i'$ and $j \neq j'$, then $M(i,j') \neq \lambda$ and $M(i',j) \neq \lambda$. Conversely if this condition is satisfied for all (i,j) then M is a solution to the oblivious routing.*

Proof: If either $M(i,j')$ or $M(i',j)$ is λ then any permutation $\pi(.)$ such that $\pi(i) = j$ and $\pi(i') = j'$ can not be satisfied. Conversely if the condition holds, then a connection from i to j can always be carried on λ without interfering with any other communication. $\square$

Define a *legal coloring* of an $n \times n$ matrix M to be an assignment of colors to the entries of M with the following property: if λ is the color of $M(i,j)$ and $M(i',j')$ for $i \neq i'$ and $j \neq j'$, then $M(i,j')$ and $M(i',j)$ are not colored with λ. The above lemma reduces the oblivious routing problem to the problem of finding a legal coloring of an $n \times n$ matrix with a minimum number of colors. We first prove that $\lceil n/2 \rceil + 2$ colors are needed and then construct an optimal solution with $\lceil n/2 \rceil + 2$ colors.

Theorem 12: *For $n \geq 6$ and $n = 4$, any legal coloring of an $n \times n$ requires at least $\lceil n/2 \rceil + 2$ colors.*

Proof: To obtain a contradiction, assume that we are given a legal coloring with $\lceil n/2 \rceil + 1$ colors. We mark each entry of the matrix with either R or C according to the following rule:

An entry $M(i,j)$ is an R-entry if its color appears more than once in row i; it is a C-entry if its color appears more than once in column j. In case its color does not appear again in both row i and column j it is marked arbitrarily.

Since the coloring is legal it follows that an R-entry cannot match the color of any other entry in its column and a C-entry cannot match the color of any other entry in its row. For each line (row or column), let $N(line)$ be the number of entries in the line marked *compatibly* with the line: $N(row)$ counts the number of R-entries in that row, and $N(column)$ counts the number of C-entries in that column. It follows that the sum of $N(line)$ over all $2n$ lines is n^2 since each entry is counted once, either in its row or in its column. Thus the average value for $N(line)$ is $n/2$.

Assume now that $n \geq 4$ is even. Since all the C-entries in a row (or all R-entries in a column) are colored with different colors, it follows that the number of colors in each line is at least $1 + n - N(line)$. Consequently, $1 + n - N(line) \leq n/2 + 1$ which implies that $N(line) \geq n/2$. However, since the average value for $N(line)$ is $n/2$, it must be that $N(line) = n/2$. Since the number of colors is $n/2 + 1$, all of the lines have the following structure: one color appears $n/2$ times and each of the other $n/2$ colors appears exactly once. We refer to the color that appears $n/2$ times as the *dominating color* of the line. For $n \geq 4$: $2 \cdot (n/2 + 1) < 2n$. Therefore, there are three lines with the same dominating color, say c. Without loss of generality, assume that two of these lines are rows. We claim that in this case c cannot appear in any entry outside these two rows – a contradiction. To see this, note that since the coloring is legal, the entries colored c in these two rows cannot share a column. Since there are n entries colored c in these two rows, for every column in the matrix, there is an entry in one of these rows colored c. However, this implies that c cannot appear anywhere else in all of these columns.

Assume now that $n \geq 7$ is odd. In this case we assume that we are given a legal coloring with $(n+3)/2$ colors. Similar arguments to the even case show that $1 + n - N(line) \leq (n+3)/2$. This implies that $N(line) \geq (n-1)/2$. However, since the average value for $N(line)$ is $n/2$, it follows that there are at least n lines with $N(line) = (n-1)/2$. Since the number of colors is $(n+3)/2$, all these lines have the following structure: one color appears $(n-1)/2$ times and each of the other $(n+1)/2$ colors appears exactly once. We refer to the color that appears $(n-1)/2$ times in such a line as the *dominating color* of the line, and to the line as a *dominated line*. For $n \geq 7$, the number of colors $(n+3)/2$ is strictly less than the number of dominated lines n. Therefore, there are two dominated lines with the same dominating color, say c. Suppose that these two lines are one row and one column. Consider the entry where this row and this column intersect. This entry cannot be colored by c. If this entry is an R-entry (respectively, a C-entry), then this row (respectively, column) has at least $(n-1)/2 + 1$ entries marked R (respectively, C), contradicting the definition of a dominated line. Thus, these two lines are either both rows or both columns. Without loss of generality assume that both are rows. The entries colored c in these two rows cannot share a column. Since there are $n-1$ entries colored c in these two rows, there is only one column where color c may color entries not in these two rows. So, if we eliminate these two rows, and the one column, we are left with an $(n-2) \times (n-1)$ matrix that is legally colored with $(n+1)/2$ colors.

We proceed to show that this is impossible.

As before, we mark the $(n-2)\cdot(n-1)$ entries with R and C. By similar arguments, we get that $N(row) \geq (n-1)/2$ and $N(column) \geq (n-3)/2$. We refer to the lines for which $N(row) = (n-1)/2$ or $N(column) = (n-3)/2$ as dominated lines. In a dominated line there must be a dominating color appearing $N(line)$ times while all the other colors appear exactly once. To lower bound the number of dominated lines note that if for all rows, $N(Row) > (n-1)/2$ and for all columns $N(column) > (n-3)/2$, the sum of $N(line)$ over all lines is at least $\frac{n+1}{2}(n-2) + \frac{n-1}{2}(n-1) = \frac{2n^2-3n-1}{2}$. However, there are only $(n-1)(n-2)$ entries in the matrix. Therefore, there are at least $\frac{2n^2-3n-1}{2} - (n-1)(n-2) = \frac{3n-5}{2}$ dominated lines.

For $n > 7$: $2\cdot(n+1)/2 < (3n-5)/2$. Therefore, there exists a color which dominates at least three dominated lines. Following the same arguments as before it can be shown that these three lines cannot be either all rows or all columns. Thus, one of these lines is a row and one is a column. We get a contradiction by examining the entry where these lines intersect, as before.

The remaining case is when $n = 7$, and the number of colors is four. Again, no color dominates three lines or one line and one column. Since $2\cdot(7+1)/2 = (3\cdot 7-5)/2$, it follows that each of the four colors dominates exactly two lines. One of them must dominate two columns because there are only $7-2=5$ rows. Moreover, this color appears in these two dominated columns and in at most one row. If we omit these three lines, we are left with a 4×4 matrix that is colored legally with the remaining three colors. This is impossible by the even case proved earlier. $\square$

We note that for $n = 2,3,4$ we need n colors to cover the matrix. The case $n = 5$ is unique since we can color a 5×5 matrix with $4 < \lceil 5/2\rceil + 2$ colors as shown in Figure 7.

$$\begin{array}{ccccc}
1 & 1 & 3 & 4 & 2 \\
3 & 4 & 1 & 1 & 2 \\
2 & 2 & 3 & 4 & 1 \\
3 & 4 & 2 & 2 & 1 \\
4 & 3 & 4 & 3 & 1
\end{array}$$

Figure 7: Optimal solution for the oblivious routing for $n = 5$ with 4 wavelengths

Now, we construct a solution using $\lceil n/2\rceil + 2$ wavelengths for $n \geq 6$. We will construct a matrix M satisfying the conditions of Lemma 11. For an even n, the idea of the construction is well demonstrated by the routing matrix presented in Figure 8. In the example $n = 12$, and the wavelengths are denoted by $0,\ldots,7$. In general, for an even n, we have the following matrix. The entries of wavelength 0 are $M_n[0,i]$, for $i = 0,\ldots,n/2-2$, $M_n[j,n/2-1]$, for $j = 1,\ldots,n/2-1$, $M_n[n/2,i]$, for $i = n/2,\ldots,n-2$, $M_n[j,n-1]$, for $j = n/2+1,...,n-1$. It is easy to see that wavelength 0 obeys the conditions of Lemma 11. Now, for $\lambda = 1,\ldots,n/2-1$, the entries of wavelength λ are given by adding λ (modulo n) to the row index and subtracting λ (modulo n) from the column index of every entry of wavelength 0. Again, it is easy to see that these wavelengths also obey the conditions of Lemma 11. The rest of the entries are filled with the two wavelengths left. The entries of wavelength $n/2$ are $M_n[i, n-i-1]$, for $i = 0, \ldots, n-1$, and the entries of wavelength $n/2+1$ are $M_n[i, n/2-i-1]$, $M_n[n/2+i, n-i-1]$, for $i = 0, \ldots, n/2-1$.

$$
\begin{array}{cccccc|cccccc}
0 & 0 & 0 & 0 & 0 & 7 & 5 & 4 & 3 & 2 & 1 & 6 \\
1 & 1 & 1 & 1 & 7 & 0 & 5 & 4 & 3 & 2 & 6 & 1 \\
2 & 2 & 2 & 7 & 1 & 0 & 5 & 4 & 3 & 6 & 2 & 2 \\
3 & 3 & 7 & 2 & 1 & 0 & 5 & 4 & 6 & 3 & 3 & 3 \\
4 & 7 & 3 & 2 & 1 & 0 & 5 & 6 & 4 & 4 & 4 & 4 \\
7 & 4 & 3 & 2 & 1 & 0 & 6 & 5 & 5 & 5 & 5 & 5 \\
- & - & - & - & - & - & - & - & - & - & - & - \\
5 & 4 & 3 & 2 & 1 & 6 & 0 & 0 & 0 & 0 & 0 & 7 \\
5 & 4 & 3 & 2 & 6 & 1 & 1 & 1 & 1 & 1 & 7 & 0 \\
5 & 4 & 3 & 6 & 2 & 2 & 2 & 2 & 2 & 7 & 1 & 0 \\
5 & 4 & 6 & 3 & 3 & 3 & 3 & 3 & 7 & 2 & 1 & 0 \\
5 & 6 & 4 & 4 & 4 & 4 & 4 & 7 & 3 & 2 & 1 & 0 \\
6 & 5 & 5 & 5 & 5 & 5 & 7 & 4 & 3 & 2 & 1 & 0
\end{array}
$$

Figure 8: Optimal solution for the oblivious routing for $n = 12$ with 8 wavelengths

Next, assume that n is odd. The matrix M_n is the matrix M_{n-1} with an additional top line and an additional last column. Since $\lceil n/2 \rceil + 2 = (n+1)/2 + 2 = (n-1)/2 + 3$, we have an extra wavelength which is denoted by $(n-1)/2 + 2$. The entries are as follows: (i) $M_n[0, j] = M_n[i, n-1] = (n-1)/2 + 2$, for $j = 0 \ldots n-2$ and $i = 1 \ldots n-1$, (ii) $M[0, n-1] = (n-1)/2$, and (iii) $M_n[i, j] = M_{n-1}[i-1, j]$, for $i = 1 \ldots n-1$ and $j = 0 \ldots n-2$.

The following theorem is implied by the above construction.

Theorem 13: *The matrix M_n is a solution to the oblivious routing using $\lceil n/2 \rceil + 2$ wavelengths.*

2.3.1. Partially Oblivious Routing

We now consider the case of partially oblivious routing. Let k be a bound on the number of wavelengths permitted to be used between any pair of transmitter receiver. In case $k = O(\log n)$ then $w = \Omega(\sqrt{n})$ wavelengths are required, $O(\sqrt{n \log n})$ are sufficient (existentially), and $\sqrt{n}\xi(n)$ are sufficient (constructively). (See Section 2.1.)

The existential upper bound for $k = o(\log n)$ can be achieved as follows. Assume as in the case of the $k = O(\log n)$ that we are looking at partitions of the receivers. If the degree is k this implies that we are looking at k different partitions. Assume that each such partition is to b blocks each of cardinality n/b. In this case the number of wavelengths is $w = kb$. To get the bound we have to find the minimum w such that the failure probability is less than 1.

Assume that $n/b \leq kb$, otherwise the failure probability is clearly 1. Let $\alpha = n/b$. Borrowing the terminology of the proof of Theorem 2, we fail if and only if there exists a set A of α vertices from S, such that $|N(A)| < \alpha$. For a fixed set $A \subset S$ and sets $A_i \subseteq B_i$ (where $a_i = |A_i|$) such that $| \cup_i A_i| = \alpha - 1$, the probability that $N(A) \subseteq \cup_i A_i$ is at most $\prod_{i=1}^{k}(\frac{a_i}{b})^\alpha$. Thus, the probability that there exist A and A_i's of cardinality α and a_i respectively such that $|N(A)| < \alpha$ is at most

$$\binom{n}{\alpha}\prod_{i=1}^{k}\binom{n/\alpha}{a_i}\left(\frac{a_i}{n/\alpha}\right)^\alpha \leq \binom{n}{\alpha}\binom{n/\alpha}{\alpha/k-1}\left(\frac{(\alpha-k)\alpha}{kn}\right)^\alpha \prod_{i=1}^{k-1}\binom{n/\alpha}{\alpha/k}\left(\frac{\alpha^2}{kn}\right)^\alpha$$
$$\leq \frac{e^{2\alpha-1}\alpha^{(2k-3)\alpha+2}}{n^{(k-2)\alpha+1}k^{(k-1)\alpha+1}} .$$

For $k = 2$, this expression is less than 1 if $\alpha \leq c \log n/\log\log n$, for some constant c. This gives $w = kn/\alpha = O(n \log\log n/\log n)$. Things look better for $k > 2$. Then, this expression is less than 1 if $\alpha \leq c \log n^{(k-2)/(2k-3)}k^{(k-1)/(2k-3)}$, for some constant c. Thus, the number of wavelengths is

$$O\left(n^{\frac{k-1}{2k-3}}k^{\frac{k-1}{2k-3}}\right) .$$

For example, for $k = 3$ it is $O(n^{2/3})$, for $k = 4$ it is $O(n^{3/5})$, and so on. As k increases the exponent of n tends from above to $1/2$. Note that for such values it is always the case that $kb = kn/\alpha \geq \alpha = n/b$.

3. Reconfigurable Optical Networks

In this section we consider reconfigurable optical networks, i.e., networks with optical switches. The number of wavelengths required to support a particular traffic set in reconfigurable networks is expected to be much smaller than in switchless networks, and is a function of the number of switches in the network. We remind the reader that the networks of this section continue to use non-reconfigurable routers which were used in Section 2.

We consider the problem of constructing a reconfigurable optical NBN. Our goal is to study the tradeoffs between the number of switches and the number of different wavelengths used in the network. As in Section 2 we differentiate between rearrangeably NBNs and wide-sense NBNs, and consider several variations of this problem. These variants arise because of different capabilities that can be attributed to the transmitters, or receivers, or the switches.

We consider two kinds of optical switches: *wavelength-selective* switches and *wavelength--independent* switches. Wavelength-selective switches, considered by Pankaj [Pan92], are more powerful than wavelength-independent switches in that they can change their state differently for different wavelengths. Wavelength-independent switches are considered in [BH92, PS93]; these switches may not be set differently for different wavelengths.

3.1. Non-Blocking Networks with Wavelength-Selective Switches

Pankaj [Pan92] considered networks with n wavelength-selective switches of fixed degree with each of the n input (and output) nodes being connected to a different switch. For this model, $\Omega(\log n)$ wavelengths are required for permutation routing [Pan92]. Pankaj also showed that rearrangeably non-blocking permutation routing can be done with $O(\log^2 n)$ wavelengths and wide-sense non-blocking permutation routing with $O(\log^3 n)$ wavelengths.

Theorem 14 proves a lower bound of $\Omega(\frac{n}{w}\log\frac{n}{w^2})$ on the number of constant degree switches required as a function of the number of nodes n and wavelengths w. Theorem 16 creates a permutation network using w wavelengths and $s = O(\frac{n}{w}\log\frac{n}{w})$ constant degree switches. In our network the input nodes are partitioned into s groups of size n/s[¶] each. Each such group is connected to an "optical combiner": a non-reconfigurable router with n/s input ports and one output port that routes any signal on any input port on to the unique output port. The output ports of the s combiners are each connected to a different switch. Similarly, the output nodes are partitioned into s groups of size n/s each. Each such group is connected to an "optical splitter": a non-reconfigurable router with one input port and n/s output ports that routes any signal coming in on the input node on to all output nodes. The input ports of the s splitters are each connected to a different switch. For the special case of n switches considered by Pankaj, we have a network that uses $O(\log n)$ wavelengths and n switches to route permutations in which each input (output) node is connected to a different switch. Thus, improving over the result obtained by Pankaj. Theorem 17 obtains a similar result for the wide-sense non-blocking network.

Theorem 14: *Permutation routing of n messages using w wavelengths requires $\Omega(\frac{n}{w}\log\frac{n}{w^2})$ wavelength-selective switches of constant degree.*

Proof: Let s denote the number of constant degree switches required, and let c be the degree of these switches. Using an argument similar to that used in [BH92], we observe that the number of switching states in the network is upper bounded by $c^{ws}w^{2n}$. This must be greater than the number of "traffic states" in the network, which is $n!$ for permutation routing. Thus we have

$$c^{ws}w^{2n} \geq n!,$$

and using Stirling's approximation this yields $s = \Omega(\frac{n}{w}\log\frac{n}{w^2})$. □

Our constructions are based on the following proposition given in [Lei92, Thm. 1.16,p.190].

Proposition 15: *Given any permutation ρ from $k\ell$ elements to $k\ell$ elements $\{x_{ij}\}_{i=1,j=1}^{k,\ell}$, ρ can be expressed as the product of three permutations ρ_1, ρ_2 and ρ_3, where ρ_1 and ρ_3 preserve the "row" index of the elements and ρ_2 preserves the "column" index. (A permutation σ preserves the "row" index if there are k permutations $\sigma_1,\ldots,\sigma_k$ from ℓ elements to ℓ elements, such that for each x_{ij}, $\sigma(x_{ij}) = x_{ij'}$, where $j' = \sigma_i(j)$. A permutation that preserves the "column" index is defined similarly.)*

[¶] To avoid cumbersome notation we assume that s divides n. Otherwise, $\lfloor\cdot\rfloor$ and $\lceil\cdot\rceil$ have to be added appropriately.

Theorem 16: *We can construct an optical reconfigurable rearrangeably NBN with w wavelengths and $O(\frac{n}{w}\log\frac{n}{w})$ wavelength-selective 4×4 switches.*

Proof: The network G we construct uses a "traditional" rearrangeably non-blocking network H for $\frac{n}{w}$ inputs and $\frac{n}{w}$ outputs, as a "black box". It is well-known that such networks using $s = O(\frac{n}{w}\log\frac{n}{w})$ 2×2 switches exist. (See, e.g. [Lei92].) Such networks can be constructed in $m = O(\log\frac{n}{w})$ layers of $\frac{n}{w}$ switches each, where the input nodes are connected to the first layer of switches and the outputs are connected to the last layer.

The switches of G correspond to the switches of H, and the first two input and output ports of each 4×4 switch of G are connected as the corresponding switch of H. Note however, that since the switches of G are wavelength-selective, we can view this portion of G as w rearrangeably NBNs superimposed, one for each wavelength. Denote a switch of G by $S_{i,j}$, where i denotes its layer index and j denotes its position within the layer. The third and fourth ports of the switches connect all the switches in the same position in all the layers cyclically. Specifically, the third (fourth) output port of $S_{i,j}$ is connected to the third (resp. fourth) input port of $S_{i+1,j}$, where the first layer is considered the successor of the last layer.

In addition, as mentioned above, the input nodes are partitioned into s groups of size n/s each. Each such group is connected to an "optical combiner", whose output is connected to a different switch. Similarly, the output nodes are partitioned into s groups of size n/s each. Each such group is connected to an "optical splitter", whose input is connected to a different switch. We partition the (input and output) nodes into n/w sets of w nodes each, where each such set corresponds to the nodes connected to switches in the same position in all the layers, and view them as w columns of size n/w each.

To route a permutation ρ in this network, we decompose ρ into $\rho_1\cdot\rho_2\cdot\rho_3$ using Proposition 15 above with $k=\frac{n}{w}$ and $\ell = w$, according to the partition of the nodes into w columns of size n/w each. Let $\rho(i,j) = (i',j')$, and let $\rho_1(i,j) = (i,j'')$, $\rho_2(i,j'') = (i',j'')$, $\rho_3(i',j'') = (i',j')$. We assign the input (i,j) the wavelength $\lambda_{j''}$, and route it as follows. Using the third ports we route (i,j) to switch $S_{1,j}$. Note that since each input in a row is assigned a different wavelength, this can be done. Then, using the rearrangeably NBN for wavelength $\lambda_{j''}$ we route (i,j) to the fourh output port of $S_{m,j'}$ Finally, using the fourth ports we route (i,j) from $S_{m,j'}$ to $S_{i',j'}$. Again, since all the outputs in a row are assigned a different wavelength, this can be done. $\square$

Theorem 17: *We can construct an optical reconfigurable wide-sense NBN with $2w-1$ wavelengths and $O(\frac{n}{w}\log\frac{n}{w})$ wavelength-selective switches.*

Proof: The idea here is similar to that of Theorem 16. The network here is the same as the one above, except for two differences: (1) We replace the rearrangeably NBN H in that construction with a wide-sense NBN with n/w inputs and n/w outputs. It is well-known that such networks using $s = O(\frac{n}{w}\log\frac{n}{w})$ switches exist [ALM90]. (2) To route a connection establishment request from input (i,j) to output (i',j') we look for a wavelength λ_k that is not currently in use at the third ports of switches in position i in all the layers, and also not in use at the fourth ports of switches in position i' in all the layers. At the time of establishing a new connection, at most $w-1$ wavelengths are used at the third ports of switches in position i (to connect at most $w-1$ other inputs) and at most $w-1$ wavelengths are used at the fourth ports of switches in position i' (to connect at most $w-1$ other outputs). Since there are a total of $2w-1$ wavelengths, we are thus guaranteeed to find a common wavelength that is not in use in all these ports. $\square$

3.2. Non-Blocking Networks with Wavelength-Independent Switches

Networks with wavelength-independent switches merit consideration since wavelength-selective switches are much harder to build than wavelength-independent switches. As for the wavelength-selective case there is a trade-off between the number of switches and the number of wavelengths used. Barry and Humblet [BH92] showed that in a wavelength-independent switch network with w wavelengths, the number of switches must be $\Omega(n \log \frac{n}{w^2})$ for both rearrangeable and wide-sense NBNs. Pieris and Sasaki [PS93] constructed such networks that use $O(n \log \frac{n}{w})$ wavelength-independent switches. Here, we show tighter upper bounds on the number of switches required in such networks by combining the arguments from Sections 2 and the previous subsection.

Theorem 18: *Given w wavelengths, there exists an optical rearrangeable NBN of size n that uses $O(n \log \frac{n \log w}{w^2})$ wavelength-independent switches. Furthermore, we can* construct *an optical rearrangeable NBN that uses $O(n \log \frac{n\xi(n)}{w^2})$ wavelength-independent switches.*

Proof: Our construction is shown in Figure 9. Again, we use Proposition 15 about decompositions of permutations. We first present an informal description: The network is constructed in three layers. The first and third layer are used to do the "row" permutations and the second layer does the "column" permutation. The first and third layer are constructed from m wavelength-independent switches with n/m inputs and n/m outputs. Each such switch is basically a "traditional" rearrangeable NBN and thus can be built from $\frac{n}{m}\log(\frac{n}{m})$ 2×2 switches. (See, e.g. [Lei92].) The middle layer is constructed from n/m copies of a non-reconfigurable network with m inputs and outputs. (The parameter m is to be chosen as a function of the number of wavelengths w as described below.) We now describe the construction more formally.

Fix $m = \frac{w^2}{\log w}$. Let F be a non-reconfigurable rearrangeable NBN for m inputs and m outputs. Such a network exists, as shown in Section 2. We will use $\frac{n}{m}$ copies of F labelled $F_1, \ldots, F_{\frac{n}{m}}$. Let H be a "traditional" rearrangeable network with $\frac{n}{m}$ inputs and outputs. We will need $2m$ copies of H labelled H_{1j} and H_{2j} for $j \in \{1, \ldots, m\}$. Let the inputs for our optical network G be labelled (i,j) for $i \in \{1, \ldots, m\}$ and $j \in \{1, \ldots, \frac{n}{m}\}$. Then, the edges of G (apart from the edges within F's and the H's) are as follows:

The input (i,j) of G is connected to the jth input of H_{1i}.

The jth output of H_{1i} is connected to the ith input of F_j.

The ith output of F_j is connected to the jth input of H_{2i}.

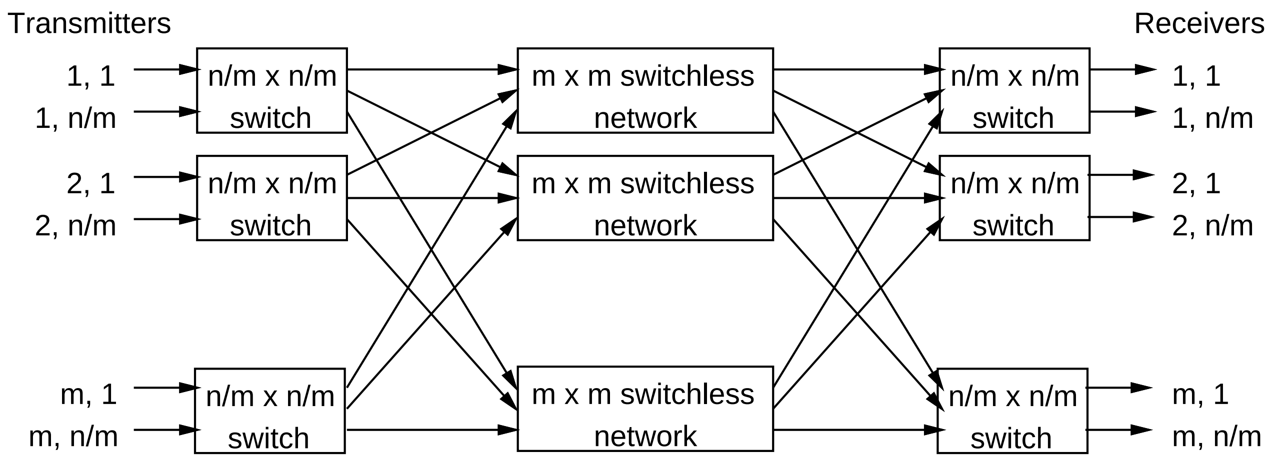

Figure 9: A permutation network using w wavelengths and $O(n \log \frac{n}{w})$ wavelength-independent switches.

The jth output of H_{2i} is connected to the (i, j)th output of G.

All edges except those within the F's can carry any of the legal wavelengths, but in the way they are used will only carry one wavelength at a time. Observe that the network uses $n \log \frac{n}{m} = n \log \frac{n \log w}{w^2}$ switches.

We now show how to route a permutation ρ. We use the row-column-row decomposition with $k = m$ and $\ell = \frac{n}{m}$ and decompose ρ into $\rho_1 \cdot \rho_2 \cdot \rho_3$. Consider a source (i, j) and let $\rho_1(i, j) = (i, j')$ and $\rho_2(i, j'') = (i', j'')$ and $\rho(i', j'') = (i', j')$. Then the path assigned to this request is from the (i, j)th source to the j''th output of H_{1i}. From there it gets routed through the network $F_{j''}$ to the i'th output node of $F_{j''}$. From there on we use the network $H_{2i'}$ to continue the path on to the sink (i', j'). Note that at most m sources use the j''th column in ρ_2. All these requests use the network $F_{j''}$ to route these pairs from their row source to their row destination. In order to do so, we pick their wavelengths as dictated by $F_{j''}$. This choice fixes the wavelength of all connections routed using the j''th column including the one with source (i, j). This determines the path and wavelength assignment of a source-sink pair. It is an easy exercise to verify that no two paths of the same wavelength overlap on any edge.

Lastly by reducing m above to $w^2/\xi(w)$, we can use the constructive non-reconfigurable rearrangeable NBN above. This yields the constructive bound claimed in the assertion. $\square$

Theorem 19: *Given w wavelengths, there exists an optical wide-sense NBN of size n that uses $O(n \log \frac{n \log w}{w^2})$ wavelength-independent switches. Furthermore, we can* construct *an optical wide-sense NBN that uses $O(n \log \frac{n\xi(n)}{w^2})$ wavelength-independent switches.*

Proof: The proof is similar to the proof of Theorem 18. The network here is almost the same as the one above, except for some minor differences: The first layer uses m n/m-input, $2n/m$-

output "traditional" wide-sense NBNs. The middle layer consists of $2n/m$ wide-sense switchless NBNs with m-inputs and m-outputs as described in Section 2. The third layer uses m $2n/m$-input n/m-output wide-sense NBNs that can route any permutation of upto n/m elements at a time. (It is easy to see that the construction in [ALM90] also yields imbalanced wide-sense NBNs that are required here. All one has to do is use a $2n/m$ to $2n/m$ NBN and throw away half the inputs or outputs as the case may be.)

Similar to the proof of Theorem 17, we can argue that any sequence of connection requests with at most one connection requested from any source or sink at any time can be routed in this network. In particular to satisfy a request from (i, j) to (i', j') we first determine the correct middle level network to use to set up this connection. For each of the networks H_{1i} and $H_{2i'}$ upto $n/m-1$ of the networks from the middle layer may already be in use. This still allows for at least one (actually at least two) networks from the middle level which neither network is using. This network can now be used to set up this connection. $\square$

Finally, we consider the special case where only the transmitters can be tuned whereas the wavelengths of the receivers are fixed (or vice-versa). For this case we have the following theorem.

Theorem 20: *Given w wavelengths, we can construct an optical wide-sense NBN of size n in which only the transmitters (receivers) are tunable that uses $\Theta(n \log \frac{n}{w})$ wavelength-independent switches of constant degree.*

Proof: We prove for the case in which only the transmitters are tunable. The other case is similar. The lower bound proof is similar to the proof of Theorem 14. Let s denote the number of constant degree switches required, and let c be the degree of these switches. Using an argument similar to that used in [BH92], we observe that the number of switching states in the network is upper bounded by $c^s w^n$. This must be greater than the number of "traffic states" in the network, which is $n!$ for permutation routing. Thus we have

$$c^s w^n \geq n!,$$

and using Stirling's approximation this yields $s = \Omega(n \log \frac{n}{w})$. The upper bound proof is similar to the proof of Theorem 19. We use a three-layer construction with the first layer using m copies of n/m to $2n/m$ "traditional" wide-sense NBNs. The middle layer construction is $2n/m$ copies of a switchless NBN but different from the one in the proof of Theorem 19. The third layer is again m copies of $2n/m$ to n/m "traditional" wide-sense NBNs.

Since in this network the receivers cannot be tuned, the destination of a request determines the wavelength of a connection. This prevents us from using the NBNs used in Theorem 19. Instead, we use a non-reconfigurable NBN that uses w wavelengths to route upto w messages with non-tunable receivers. (Construction of such a network is easy as shown in Figure 4.) This forces $m = w$ and shows that the number of switches is at most $O(n \log \frac{n}{w})$.

To route a message from input (i, j) to output (i', j') the message is allotted to the kth non-reconfigurable network in the middle layer if no message with source from ith row or destined to the i'th row is currently using the kth network. Once again since the number of such networks is twice the number of elements per row, there must be an unused network which can be used to achieve this permutation. □

3.3. Bounding Number of Wavelengths via Congestion and Dilation

In this section, we give bounds on the number of wavelengths required to route a set of messages on optical networks with n nodes, each having a wavelength-selective switch. Our bounds here relate the number of wavelengths to two classical parameters associated with routing: the congestion and the dilation.

Given a graph G with n nodes and m edges, suppose some messages can be routed in it such that the maximum congestion (i.e., number of messages using any edge) is c and the maximum dilation (i.e., the maximum path length from any source to any sink) is d. Clearly at least c wavelengths are required in order to realize the routing. Now construct a new graph G_p with each path in G being a node in G_p with an edge between two nodes in G_p if the corresponding paths in G overlap on any edge in G. Then, the problem of assigning wavelengths to paths in G reduces to that of coloring nodes in G_p. Since the maximum degree of a node in G_p is $(c-1)d$, $(c-1)d+1$ wavelengths are sufficient to achieve this routing. Our first observation shows that if the dilation is sufficiently large then it is possible to beat this bound. In particular, if $d > \sqrt{m}$ then $2c\sqrt{m}$ wavelengths suffice to route the messages. Next we show that this bound is optimal up to a constant factor, in that there exist graphs and message requests for which $\Omega(c \min\{d, \sqrt{m}\})$ wavelengths are required to route the given messages.

Lemma 21: *For any graph G, and any set of routing requests with congestion of c, $2c\sqrt{m}$ wavelengths are sufficient to achieve the given routing.*

Proof: The number of paths of length at least $\sqrt{m}$ is at most $c\sqrt{m}$. Give each such path its own dedicated wavelength. Each of the remaining paths (of length less than $\sqrt{m}$), conflicts with fewer than $c\sqrt{m}$ paths, so again can be given a wavelength without conflicting by a greedy coloring of the interference graph of the paths of length less than $\sqrt{m}$. □

Theorem 22: *There exist graphs and message routing requests which can be routed with congestion c and dilation d, but require $w = \Omega(c \min\{d, \sqrt{m}\})$ wavelengths under any routing.*

Proof: Define the graph G with n transmitters and n receivers as depicted in Figure 10. In addition to the transmitters and receivers, G has n columns, each column i consists of $2n-1$ nodes: $a_{i,j}$, $j = 1, \ldots, n$ and $b_{i,j}$, $j = 1, \ldots, n-1$. The edges are as follows: For $i = 1, \ldots, n$, t_i is connected to $a_{1,i}$ and r_i is connected to $a_{n,i}$. For $j = 1, \ldots, n-1$ and $i = 1, \ldots, n$, $b_{i,j}$ is connected to $a_{i,j}$ and $a_{i,j+1}$. For $j = 1, \ldots, n$ and $i = 3, 5, \ldots, n$, $a_{i,j}$ is connected to $a_{i-1,j}$. For $j = 1, \ldots, n-1$ and $i = 2, 4, \ldots, n$, $b_{i,j}$ is connected to $b_{i-1,j}$. Note that the number of edges is $O(n^2)$.

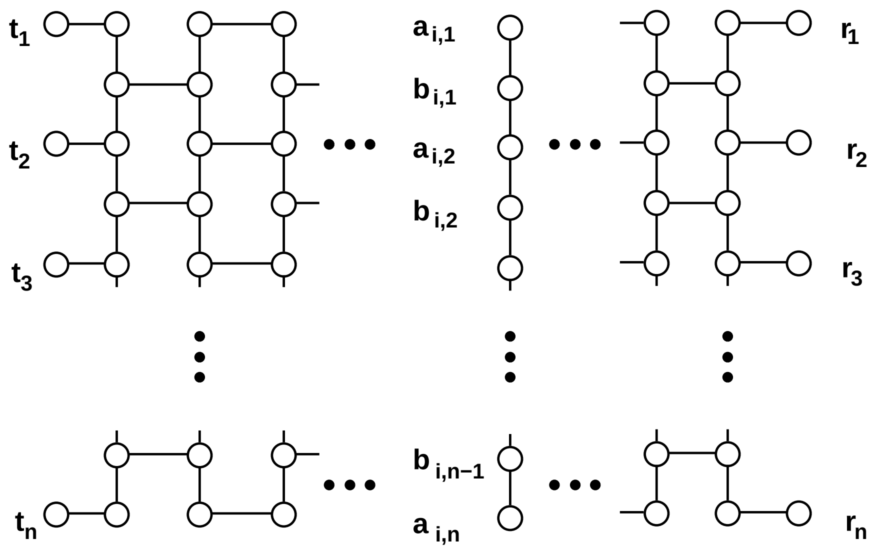

Figure 10: The network example used to prove Theorem 22.

There are n requests. The transmitter and the receiver of the ith connection are nodes t_i and r_{n-i+1} respectively. It is not hard to verify that for any choice of n paths satisfying these n connection, any two paths intersect in at least one edge. Consequently each path needs a unique wavelength, which implies that the number of wavelengths is n.

It is possible to construct n paths such that the congestion is constant and the dilation is $O(n)$. This is done as follows. For $1 \le i < n/2$, the path from t_i to r_{n-i+1} is composed of three parts: (i) to $a_{1,i}$ and then "straight" to $a_{i,i}$ (using nodes $b_{j,i}$ and $a_{j,i}$), (ii) down to $b_{i,n-i+1}$, (iii) "straight" to $a_{n,n-i+1}$ and r_{n-i+1} (using nodes $b_{j,n-i+1}$ and $a_{j,n-i+1}$). The paths for $n/2 < i \le n$ and for $i = n/2$ (in case n is even) are symmetric. Since in this graph $n = O(\sqrt{m})$, it follows that $w = \Omega(\min\{d, \sqrt{m}\})$.

By letting replacing each original transmitter by c transmitters all connected in to the original transmitter, and the same for the original receivers, we can force any value for the congestion c. The theorem follows since the number of wavelengths in this case is $O(cn)$. □

4. Open Problems

There are several unresolved problems related to the models in this paper; some of them are listed below:

1. We do not have tight bounds for switchless non-oblivious networks. The same is also true for the switchless, partially oblivious networks that allow at most k wavelengths on any edge.

2. Our constructive algorithms for the wide-sense, non-oblivious networks take exponential time

in many cases; it would be useful to obtain polynomial time constructive algorithms for these cases.

3. We provided an algorithm to convert any network with m edges that routes messages with congestion c, dilation d, into a network that uses $O(c(\min(d, \sqrt{m})))$ wavelengths, and we also gave a network and a message pattern for which this bound is optimal up to a constant factor However, a much more interesting and practically useful question is getting a good bound on the number of wavelengths required for a given network and a given message pattern. Here, we have no results and getting even an approximate bound on the number of wavelengths would be very interesting.

4. One research topic not studied in this paper is that of strict-sense non-blocking networks. A strict-sense non-blocking network is one that allows a new connection to be always routed through irrespective of how the previous connections were routed.

5. Another topic not studied here is the use of wavelength converters. A wavelength converter can convert a signal from one wavelength to another. Clearly all the upper bounds in the paper hold for networks with wavelength converters. We can also show that the lower bounds hold for networks with *static* wavelength converters. With dynamic converters, we are adding states to the network and hence the lower bounds may not apply.

   A related model not studied here is that of networks that are not all-optical. In these networks, a connection need not be carried on a single wavelength all the way to its destination; it could be carried on one wavelength to an intermediate node, where it is received and switched electronically onto another wavelength enroute to its destination. In this case, it is also possible to multiplex several connections on to a single wavelength, allowing packet switching.

6. The area of fault tolerance in optical networks is an open area for research. Also, there is the set of problems of dynamically maintaining topology of optical networks (especially when the links are created and/or destroyed), and maintaining information regarding link utilization, congestion, etc.

Acknowledgments

We wish to thank the anonymous referees for numerous comments which have hopefully improved the clarity of this paper.

References

[ALM90] S. Arora, T. Leighton, and B. Maggs. On-line algorithms for path selection in a non-blocking network. In *Proc. of the ACM Symp. on Theory of Computing*, pages 149–158, 1990.

[Ben62] V.E. Beneš. Heuristic remarks and mathematical problems regarding the theory of switching systems. *Bell System Tech. Journal*, 41:1201–1247, 1962.

[BH92] R. A. Barry and P. A. Humblet. Bounds on the number of wavelengths needed in WDM networks. In *LEOS'92 Summer Topical Mtg. Digest*, pages 21–22, 1992.

[BH93a] R. A. Barry and P. A. Humblet. On the number of wavelengths and switches in all-optical networks. *IEEE Trans. on Communications, vol. 42, no.2/3/4, pp. 583–591, 1994.*

[BH93b] R. A. Barry and P. A. Humblet. An all-optical non-blocking $M \times M$ switchless connector with $O(\sqrt{M \log M})$ wavelengths and without wavelength changers. *Electron Lett.*, 29:1252-1254, July 1993.

[BRSU93] A. Borodin, P. Raghavan, B. Schieber, and E. Upfal. How much can hardware help routing? In *Proc. of the 25th ACM Symp. on Theory of Computing*, San Diego, CA, May 1993.

[CNW90] N. K. Cheung, K. Nosu and G. Winzer (eds.) *IEEE JSAC: Special Issue on Dense WDM Networks.* volume 8, Aug. 1990.

[FFP88] P. Feldman, J. Friedman, and N. Pippenger. Wide-sense nonblocking networks. *SIAM Journal on Discrete Math.*, 1(2):158–173, May 1988.

[Gre92] P. E. Green. *Fiber-Optic Networks.* Prentice Hall, 1992.

[Hal35] P. Hall. On the representatives of subsets. *Journal of London Mathematical Society*, 10(1):26–30, 1935.

[Lei92] F.T. Leighton. *Introduction to parallel algorithms and architectures: arrays, trees, hypercubes.* Morgan Kaufmann Publishers, San Mateo, CA, 1992.

[Pan92] R.K. Pankaj. *Architectures for linear lightwave networks.* PhD thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA, September 1992.

[PS93] G. R. Pieris and G. H. Sasaki. A linear lightwave Beneš network. *IEEE/ACM Trans. on Networking*, 1:441-445, Aug. 1993.

[Ram93] R. Ramaswami. Multi-wavelength lightwave networks for computer communication. *IEEE Communications Magazine*, 31(2):78–88, Feb. 1993.

[WZ93] A. Wigderson and D. Zuckerman. Expanders that beat the eigenvalue bound: explicit construction and applications. In *Proc. of the 25th ACM Symp. on Theory of Computing*, San Diego, CA, May 1993.