Strictly Non-blocking WDM Cross-connects for Heterogeneous Networks

April Rasala *
MIT Laboratory for Computer Science
Cambridge, MA 02139
arasala@theory.lcs.mit.edu

Gordon Wilfong
Bell Labs, Lucent Technologies
Murray Hill, NJ 07974
gtw@research.bell-labs.com

Abstract

A wavelength division multiplexed (WDM) network is a network of optical fibers and components that supports multiple wavelength channels along each fiber. A heterogeneous WDM network is one that is composed of sub-networks each supporting possibly different sets of wavelength channels. Thus the number of wavelength channels, n_1, on the input fibers to a cross-connect connecting one such sub-network to another will possibly differ from the number of wavelength channels, n_2, on its output fibers. In order to balance the number of channels coming into and going out of such a cross-connect it is possible then that the number of input fibers, k_1, will differ from the number of output fibers, k_2. We study the design of such cross-connects. In particular we consider *strictly non-blocking* WDM cross-connects where a cross-connect C is strictly non-blocking if it is always capable of routing any unused wavelength channel on any fiber entering C onto any unused wavelength channel on any fiber leaving C regardless of any other routes already in use within C. Any such WDM cross-connect will need some number of *wavelength interchangers*, expensive optical components that translate incoming wavelength channels onto different outgoing wavelength channels. We provide a simple method for constructing a $k_1 \times k_2$ strictly non-blocking WDM cross-connect having at most $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers for any $k_1 \leq k_2$. A proof is then given showing that these constructions are optimal in that any such WDM cross-connect for a heterogeneous network must have at least $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers.

*Supported by a Lucent Technologies GRPW fellowship.

1 Introduction

Wavelength division multiplexing (WDM) refers to the technology that allows multiple signals to pass simultaneously along an optical fiber where each such signal is encoded on a distinct wavelength of light over the fiber [SB99]. Thus each fiber can be thought of as supporting some number of *wavelength channels*. Complex WDM networks require WDM cross-connects that permit switching at the wavelength granularity level. That is, the function of a WDM cross-connect C is to allow signals on wavelength channels on fibers coming into C to be routed onto wavelength channels on fibers going out of C. *Wavelength interchangers*, optical devices that can translate incoming wavelength channels onto different outgoing wavelength channels, will be required in such cross-connects. Currently and in the foreseeable future, the cost of wavelength interchangers will dominate the cost of the other components in WDM cross-connects. Thus our goal is to study designs for WDM cross-connects that use as few wavelength interchangers as possible.

Also, since optical buffering is difficult and expensive, it is desirable that WDM cross-connects be strictly non-blocking, that is they should always be able to route any unused input wavelength channel onto any unused output wavelength channel no matter what other routes are currently being used. Thus in a strictly non-blocking cross-connect, no connections currently established need ever be disrupted and hence buffered, in order to satisfy new demands for connections.

There is a vast literature concerned with the problem of optimizing the design of cross-connects in the case where each input and output line carries only a single channel [Ben35, Sha50, PV76, BP74, Pip82]. Such cross-connects are called *space domain cross-connects*. In this case, since a demand never changes wavelengths and therefore no wavelength interchangers are needed in the cross-connect, the optimization criterion can be thought of as minimizing the number of edges in the graph representing the topology of the cross-connect.

There has also been some work on the problem of optimally placing wavelength interchangers in WDM networks to achieve various routing properties without having to increase the number of wavelengths available. For instance, [WW98, KK99] considered the problem of minimizing the number of wavelength interchangers needed so that the resulting network would have the same routing capabilities as if wavelength interchangers were placed at every node. Others have studied problems dealing with the possibility of using less powerful wavelength interchangers (e.g. those with the ability to only swap two wavelength channels leaving the others fixed) and asked how using such weaker wavelength interchangers affects the number of wavelength interchangers needed to achieve certain properties [RS97]. Another line of research has been the study of the optimal placement of wavelength interchangers for various network topologies and traffic models to minimize blocking probabilities [SAS98, TS99].

The problem that we address, namely the design of strictly non-blocking WDM cross-connects, has been studied in a few recent papers for the case where the number of input fibers and the number of output fibers was k and the wavelength channels on both the input and output fibers were the same. For instance, a number of WDM cross-connect designs with various non-blocking properties were considered in [WMDZ99] and in particular one architecture was presented for a $k \times k$ strictly non-blocking WDM cross-connect that required $k \log k$ wavelength interchangers. Later it was shown that $2k-1$ wavelength interchangers are necessary and sufficient for such strictly non-blocking WDM cross-connects [RW00].

A network may be *heterogeneous* in that it may be composed of various sub-networks that are produced by different manufacturers. Since the set of wavelength channels supported in a sub-network can depend on the particular manufacturer, there is a need for WDM cross-connects between sub-networks that support different wavelength channels. Such cross-connects may require the number of input and output fibers to be different. For instance, if the number of input wavelength channels is larger than the number of output wavelength channels, it is likely that the number of output fibers would have to be larger than the number of input fibers in order to handle the traffic entering the network on the larger number of input wavelengths.

This paper considers WDM cross-connects for networks in which the number of input and output fibers and input and output wavelengths may differ. We begin by describing constructions for $k_1 \times k_2$ WDM cross-connects. We can assume, without loss of generality, that $k_1 \leq k_2$. If there are n_1 input wavelengths and n_2 output wavelengths, then our designs use $min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. These are shown to be strictly non-blocking in a straightforward manner. In the case where $k_2 \geq n_1 k_1$ it is easy to show that the number of wavelength interchangers needed matches that of our construction. When $k_2 < n_1 k_1$, showing that our designs are optimal is much more involved. We use an iterative technique to show that $\min(k_1, n_1 k_1 - k_2) + k_2 - 1$ wavelength interchangers are required and so our construction is optimal if $k_2 \leq (n_1 - 1)k_1$. However if $(n_1 - 1)k_1 < k_2 < n_1 k_1$ then this result only says that $n_1 k_1 - 1$ wavelength interchangers are required while our construction uses $n_1 k_1$. We then provide a more complicated procedure to prove that in fact our construction is optimal in this last case.

2 Definitions

A WDM cross-connect C is made up of fibers and various optical components. Fibers are directed in the sense that signals must only travel in one specified direction. A fiber f is called an *input fiber* if signals come into C via f or an *output fiber* if signals leave C via f. Throughout the paper we will be considering $k_1 \times k_2$ WDM cross-connects where $I = \{I_1, I_2, \ldots, I_{k_1}\}$ and $O = \{O_1, O_2, \ldots, O_{k_2}\}$ are the set of k_1 input fibers and k_2 output fibers, respectively. Without loss of generality, we will assume that $k_1 \leq k_2$. The *input wavelengths*, $\lambda_1, \lambda_2, \ldots, \lambda_{n_1}$, are the wavelengths available on each input fiber. Any fiber that is capable of carrying signals on all the input wavelengths is called an *input side fiber*. Similarly the *output wavelengths*, $\gamma_1, \gamma_2, \ldots, \gamma_{n_2}$, are the wavelengths available on each output fiber and any fiber capable of carrying signals on all the output wavelengths is called an *output side fiber*.

One type of optical component that can be found in a WDM cross-connect is a *wavelength interchanger*. A wavelength interchanger, WI, has one input side fiber directed into it and one output side fiber directed out of it. For any input wavelength λ_i on the fiber directed into WI and any output wavelength γ_j, WI can move the signal coming in with wavelength λ_i onto its outgoing fiber using wavelength γ_j so long as WI has not moved any other incoming signal onto wavelength γ_j. Other types of optical components in a cross-connect are input side and output side *switches*. An $n \times q$ input side switch s has n fibers directed into it and q fibers directed out of it. For any input wavelength λ_i, a signal using wavelength λ_i on any fiber into s can be switched by s onto any fiber out of s using wavelength λ_i with the restriction that at most one signal using λ_i is switched onto any given outgoing fiber. An output side switch is defined analogously for the output wavelengths.

Given such components a WDM cross-connect must have the form shown in Figure 1 and will be called a *WDM split cross-connect.* The part of the cross-connect labeled F_1, called the *input side of the cross-connect*, consists entirely of input side fibers and input side switches. The *output side of the cross-connect*, F_2, contains only output side fibers and output side switches. The two sides are connected together by the set of wavelength interchangers, $WI_1, \ldots, WI_m$.

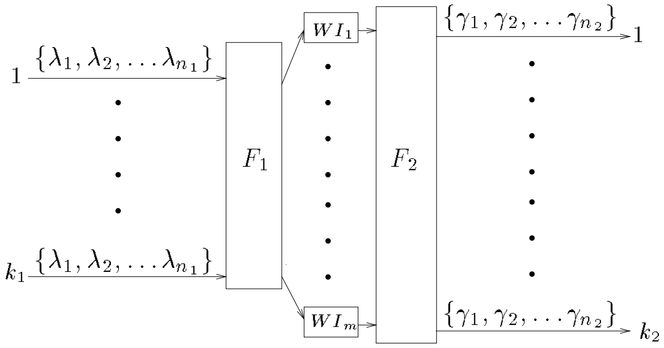

Figure 1: A WDM split cross-cross connect with k_1 input fibers and k_2 output fibers.

A *demand* d is defined as a 4-tuple $d = (I_a, \lambda_j, O_b, \gamma_l)$ meaning that there is a request to send the signal coming in on wavelength λ_j on input fiber I_a onto the output fiber O_b using wavelength γ_l. This requires finding a path in C from input fiber I_a to output fiber O_b. We call such a path a *route for demand d*. Notice that a route r for a demand $d = (I_a, \lambda_j, O_b, \gamma_l)$ consists of two parts: (1) a path within the input side of C, called the *input path*, from I_a to some wavelength interchanger WI_h and (2) a path within the output side of C, called the *output path*, from WI_h to O_b.

A *valid set of demands* is a set of demands such that no two demands have the same input fibers and input wavelengths and no two demands have the same output fibers and output wavelengths. A *routing* $R = \{r_1, r_2, \ldots, r_t\}$ for a set of demands $D = \{d_1, d_2, \ldots, d_t\}$ is a set of routes where each r_i is a route for demand d_i. We say that R is a *valid routing* if the input paths of the routes for all demands with the same input wavelength are pairwise edge (fiber) disjoint and the output paths of the routes for all demands with the same output wavelength are pairwise edge disjoint.

Suppose that D is a valid set of demands. A demand $d \notin D$ is a *valid demand* with respect to D if the demand set $D' = \{d\} \bigcup D$ is a valid set of demands. If R is a valid routing for D, then we say that r is a *valid route for d* (with respect to R) if $R' = \{r\} \bigcup R$ is a valid routing of D'. If such a valid route r exists then we say that C can *satisfy* demand d with respect to the already established routes in R. [1] If r is the route for d used by C to satisfy d then we say that the fibers and devices along r *service d*. An input wavelength λ_i (or output wavelength γ_i) is said to be *available* at a wavelength interchanger WI_j if there is no demand serviced by WI_j with input wavelength λ_i (or output wavelength γ_i).

A WDM cross-connect C is *strictly non-blocking* if for any valid routing R of a valid set of demands D and any demand $d \notin D$ such that d is a valid demand with respect to D, there is a valid route r of d with respect to R.

As mentioned earlier, a cross-connect in a network in which there is only one channel is called a *space domain* cross-connect. Space domain cross-connects are said to be *pathwise strictly non-blocking* if a connection can be made from any unused input fiber to any unused output fiber without changing the routing of any previously established connections.

Consider a WDM split cross-connect C as shown in Figure 1. For F_1, the input side of C, we define F_1^S to be the $k_1 \times m$ space domain cross-connect that has the same topology as F_1. Similarly, we define F_2^S to be the $m \times k_2$ space domain cross-connect that has the same topology as F_2. Now consider some set of valid demands D that are routed using some valid routing R in C. Let d be a valid demand with respect to D where the input wavelength of d is λ_i. Then the only demands in D that can block d from having an input path to some wavelength interchanger are those demands also with input wavelength λ_i. In other words, since demands with different input wavelengths can share input fibers without interfering with each other, the input side, F_1, is logically equivalent to n_1 copies, one for each input wavelength, of the space domain cross-connect F_1^S that has the same topology as F_1. Similarly, F_2 is logically equivalent to n_2 copies of the space domain cross-connect F_2^S. We exploit these equivalences in the design of strictly non-blocking WDM cross-connects in Section 3.

3 Constructions for strictly non-blocking WDM cross-connects

We will consider two constructions of $k_1 \times k_2$ WDM cross-connects and show that they are both strictly non-blocking.

In what follows, we will be interested in WDM split cross-connects whose input and output sides, F_1 and F_2 respectively, are such that F_1^S and/or F_2^S are pathwise strictly non-blocking. Recall that a space domain cross-connect such as F_1^S or F_2^S is pathwise strictly non-blocking if for any set of existing edge disjoint paths from inputs to outputs, there is always a path from any unused input to any unused output that is edge disjoint

[1] Note that throughout the paper we will often speak of a cross-connect as being able to satisfy a demand without explicitly saying "with respect to a routing R" since R will be obvious from context.

from each of the already existing paths. It should be noted that there are many known designs for pathwise strictly non-blocking space domain cross-connects (for example, see [Can72]).

The first class of cross-connects that we consider are constructed in such a way that there is essentially one dedicated wavelength interchanger for each possible input side of a demand. Since the number of available input wavelengths is n_1 and the number of input fibers is k_1, the number of wavelength interchangers in such cross-connects is $n_1 k_1$. Note that the number of wavelength interchangers in a dedicated WDM cross-connect is the same as the largest number of demands in any valid demand set. We define C to be a *dedicated WDM cross-connect* if C is a WDM split cross-connect where F_2, the output side of C, is such that F_2^S is an $n_1 k_1 \times k_2$ pathwise strictly non-blocking space domain cross-connect. Furthermore, the topology of F_1, the input side of C, is such that each wavelength interchanger is reachable from exactly one input fiber, each input fiber can reach exactly n_1 wavelength interchangers and the path between any input fiber and any wavelength interchanger consists of the input fiber to a $1 \times n_1$ input side switch followed by a single dedicated fiber to the wavelength interchanger. Figure 2 illustrates the design of a dedicated WDM cross-connect.

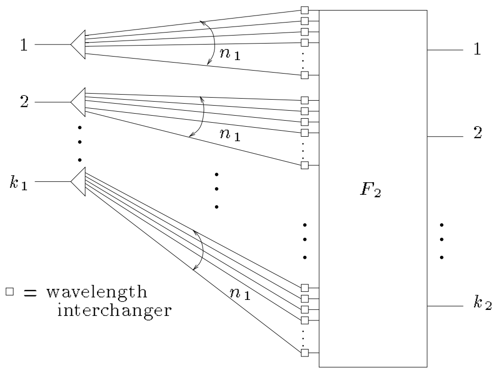

Figure 2: A dedicated WDM cross-cross connect that has n_1 wavelength interchangers *dedicated* for each of the k_1 input fibers.

Lemma 3.1 *Let C be a $k_1 \times k_2$ dedicated WDM cross-connect with n_1 input wavelengths where the number of wavelength interchangers in C is $n_1 k_1$. Then C is strictly non-blocking.*

By Lemma 3.1, any dedicated WDM cross-connect with $n_1 k_1$ wavelength interchangers is strictly non-blocking but in some cases it may be using many more wavelength interchangers than necessary. For instance, if $k_1 = k_2 = k$ then we know from [RW00] that there exists a strictly non-blocking WDM cross-connect with only $2k - 1$ wavelength interchangers and if n_1 is much larger than 2, then $n_1 k >> 2k - 1$. Therefore we consider a second class of strictly non-blocking WDM cross-connects that require only $k_1 + k_2 - 1$ wavelength interchangers.

Lemma 3.2 *Let C be a $k_1 \times k_2$ WDM split cross-connect with $k_1 + k_2 - 1$ wavelength interchangers. If the input side, F_1, and the output side, F_2, of C are such that F_1^S and F_2^S are pathwise strictly non-blocking, then C is a strictly non-blocking WDM cross-connect.*

Proof: To show that C is strictly non-blocking we need to show that for any valid routing R of demand set D and any demand $d = (a, \lambda_i, b, \gamma_j)$ that is valid with respect to D there is a valid routing r of d with respect to R. Since there are never more than k_1 demands that use the same input wavelength and never more than k_2 that use the same output wavelength in any set of valid demands, there can be at most $k_1 - 1$ and $k_2 - 1$ demands in D that use input wavelength λ_i and output wavelength γ_j respectively. Therefore there is at least $k_1 + k_2 - 1 - (k_1 + k_2 - 2) = 1$ wavelength interchanger, WI_x, that can service a demand with input wavelength λ_i and output wavelength γ_j.

Since F_1^S and F_2^S are both pathwise strictly non-blocking, there is a path from input fiber a to WI_x that shares no edge with the input side of any routing in R with input wavelength λ_i and a path from WI_x to output fiber b that shares no edge with the output side of any routing in R with output wavelength γ_j. Therefore there is a valid routing r of d. ■

Notice that as long as $k_2 \le (n_1 - 1)k_1$, then $k_1 + k_2 - 1 < n_1 k_1$ and therefore the construction in Lemma 3.2 uses fewer wavelength interchangers than the construction of a dedicated WDM cross-connect described earlier. If $k_2 > (n_1 - 1)k_1$, then a $k_1 \times k_2$ dedicated WDM cross-connect with $n_1 k_1$ dedicated wavelength interchangers is sufficient. Together then, Lemma 3.1 and Lemma 3.2 allow us to conclude the following.

Theorem 3.3 *There exists a $k_1 \times k_2$ strictly non-blocking WDM cross-connect supporting n_1 input wavelengths with $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers.*

4 Lower bounds

We begin by considering the case when $k_2 \ge n_1 k_1$ where we saw in the previous section that there exists a strictly non-blocking WDM cross-connect that uses $n_1 k_1$ wavelength interchangers.

Theorem 4.1 *If C is a strictly non-blocking $k_1 \times k_2$ WDM cross-connect with $k_2 \ge n_1 k_1$, then C must have at least $n_1 k_1$ wavelength interchangers.*

Proof: If $k_2 \geq n_1 k_1$, then a valid demand set could include $n_1 k_1$ demands that all use the same output wavelength. Since no wavelength interchanger can service more than one of these demands, any $k_1 \times k_2$ WDM cross-connect will need at least $n_1 k_1$ wavelength interchangers if it is strictly non-blocking. ∎

If $k_2 < n_1 k_1$, it is easy to see by the argument in the proof of Theorem 4.1 that any strictly non-blocking WDM cross-connect must have at least k_2 wavelength interchangers. In order to show that in this case, any strictly non-blocking WDM split cross-connect actually requires $\min(k_1+k_2-1, n_1k_1)$ wavelength interchangers, we begin by first showing that $\min(k_1, n_1k_1-k_2)+k_2-1$ wavelength interchangers are needed. To do this we consider any strictly non-blocking WDM split cross-connect C with fewer than $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers. For such a cross-connect we will create a demand set D and a routing R of D that routes D in such a way that C must use each wavelength interchanger to service a demand with either input wavelength λ_1 or output wavelength γ_2. Given this set of demands D and routing R for D we will then show that one additional valid demand $d \notin D$ exists with input wavelength λ_1 and output wavelength γ_2. Since all the wavelength interchangers in the cross-connect will already be servicing a demand that either uses this new demand's input wavelength or its output wavelength, the cross-connect will not have a wavelength interchanger to service the new demand. This implies that any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect must have at least $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers.

Consider a strictly non-blocking $k_1 \times k_2$ WDM cross-connect C where $k_2 < n_1k_1$. Let $z = \min(k_1, n_1k_1-k_2)$, $\mathcal{I}^A = \{a_1, a_2 \dots a_z\}$ be a subset of I consisting of z input fibers, $\mathcal{O}^A = \{b_1, b_2, \dots b_z\}$ be a subset of z output fibers in O and $\mathcal{O}^H = \{h_1, h_2 \dots h_{k_2-z}\}$ be the remaining output fibers. Given this partition of the input and output fibers we say that a valid demand set D is *standard* if and only if

1. $|D| = k_2 + z$,
2. $2z$ of the demands in D are demands with an input fiber in $\mathcal{I}^A$, an output fiber in $\mathcal{O}^A$, input wavelength either λ_1 or λ_2 and output wavelength either γ_1 or γ_2 and
3. the other $k_2 - z$ demands in D have output wavelength γ_2 and an output fiber in $\mathcal{O}^H$.

We will be constructing various standard demand sets D_i in what follows. For such a standard set of demands, D_i, define D_i^A to be the subset of demands in D_i with an output fiber in $\mathcal{O}^A$. The subset of demands in D_i with an output fiber in $\mathcal{O}^H$ will remain fixed for all i and we denote this set by D^H.

Suppose we have a routing R_i for a standard set of demands D_i. We will be successively transforming standard demand sets into new standard demand sets by creating new demands with input wavelength λ_1 and output wavelength γ_2 with the goal of forcing these new demands to use wavelength interchangers that are so far unaccounted for. To do this we can partition the set of wavelength interchangers in C into three sets depending on the types of demands that they service. The routes for the fixed demands in D^H will also be unchanging and the fixed set of wavelength interchangers that service these routes is denoted by W^H. The intuition is that these wavelength interchangers will always be "held" by these routes and prevent other demands with output wavelength γ_2 from being routed through them. Let W_i^B be the set of wavelength interchangers that are not in W^H yet service a demand with either input wavelength λ_1 and/or output wavelength γ_2. These wavelength interchangers are thought of as "blocking" demands with input wavelength λ_1 and output wavelength γ_2 from being routed through them. Let W_i^F be the set of all other wavelength interchangers. The wavelength interchangers in W_i^F are thought to be "free" to service demands with input wavelength λ_1 and output wavelength γ_2.

We define a valid routing R_i of a standard set of demands D_i to be *standard* if

1. each wavelength interchanger in W^H services only one demand and
2. the set of output paths of the routes for demands in D^H are edge disjoint from all output paths of the routes for demands in D_i^A.

Suppose C is a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect where $k_1 \leq k_2 < n_1k_1$. Then we show that a standard set of demands D_0 and a standard routing R_0 for D_0 exist within C. First create $z = \min(k_1, n_1k_1 - k_2)$ demands of the form $d_{i1} = (a_i, \lambda_1, b_i, \gamma_1)$ for $1 \leq i \leq z$. Clearly since C is strictly non-blocking, C must be able to satisfy these z demands. Notice also that these demands must be routed on edge disjoint paths because they all have the same input and output wavelengths. For $1 \leq i \leq z$, create a demand $d_{i2} = (a_i, \lambda_2, b_i, \gamma_2)$ and route along the same path as d_{i1}. Now create $k_2 - z$ demands so that each output fiber $h_j \in \mathcal{O}^H$ has a demand with output wavelength γ_2. By assumption, k_2 is small enough relative to k_1 and n_1 that there are enough available input wavelengths so that there are at least $k_2 - z$ such valid demands. Clearly the set of demands that use output wavelength γ_2 must all have edge disjoint output paths.

Since all other demands in D_0^A are routed along a path that also has a demand in D_0^A with output wavelength γ_2, all demands in D_0^A are routed along output paths that are edge disjoint with respect to all output paths of routes for demands in D^H. Therefore D_0 is a standard set of demands and the routing described is a standard routing R_0 for D_0.

Lemma 4.2 *If C is a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect, D_i is a standard set of demands and R_i is a standard routing of the demands in D_i that uses $\min(k_1, n_1k_1-k_2)+k_2-g$ wavelength interchangers, where $g > 0$, then g wavelength interchangers in W_i^B will each service two demands both of whose input wavelengths are λ_1 and λ_2 and whose output wavelengths are γ_1 and γ_2.*

Proof: By definition, any standard routing R_i of D_i routes at most one demand through any wavelength interchanger in W^H and therefore that demand must be a demand whose output fiber is in $\mathcal{O}^H$. There are k_2-z of these demands and therefore $|W^H| = k_2 - z$. Since the total number of wavelength interchangers used is $\min(k_1, n_1k_1-k_2)+k_2-g = z+k_2-g$ and $|W^H| = k_2-z$, it must be the case that the number of wavelength interchangers in W_i^B and W_i^F is $2z-g$. The $2z$ demands of D_i^A must be routed through wavelength interchangers in W_i^B or W_i^F since R_i is a standard routing. These $2z$ demands in D_i^A use only input wavelengths λ_1 and λ_2 and output wavelengths γ_1 and γ_2. As a result no wavelength interchanger can service more than two of these demands. This implies that g wavelength interchangers in W_i^B or W_i^F must service two demands. Any such wavelength interchanger that services two of these demands must service a demand with input wavelength λ_1 and a demand with output wavelength γ_2 and therefore is by definition in W_i^B. ■

Given the standard set of demands D_0 and the routing R_0 of D_0 we now present two manipulations that together can be used to iteratively change the set of demands on C so that eventually we arrive at a standard set of demands D_t and a standard routing R_t of D_t such that every wavelength interchanger in C is servicing a demand with either input wavelength λ_1 or output wavelength γ_2. Notice that this is equivalent to showing that if C has no more than $\min(k_1, n_1k_1 - k_2) + k_2 - g$ wavelength interchangers for some $g > 0$ then a standard set of demands D_i and a standard routing R_i of D_i exists such that $|W^H| = k_2 - z$, $|W_i^B| = 2z - g$ and $|W_i^F| = 0$.

Let $WI_j \in W_i^B$ be a wavelength interchanger that services exactly two demands d_1 and d_2. Suppose these demands have the form $d_1 = (a_1, \lambda_1, b_2, \gamma_2)$ and $d_2 = (a_2, \lambda_2, b_1, \gamma_1)$. The operation **Uncross**(WI_j) is defined to have the effect of changing these two demands to be $(a_1, \lambda_1, b_1, \gamma_1)$ and $(a_2, \lambda_2, b_2, \gamma_2)$ and routing these new demands exactly as d_1 and d_2 were routed while keeping all other demands and routes unchanged. Notice that the only real change made by **Uncross**(WI_j) is to swap the output paths of the routes of d_1 and d_2. On the other hand, if the demands have the form $d_1 = (a_1, \lambda_1, b_2, \gamma_1)$ and $d_2 = (a_2, \lambda_2, b_1, \gamma_2)$, then **Uncross**$(WI_j)$ has no effect.

Next we define the operation **Block**$(C, (D_i, R_i))$ where C is a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect, D_i is a standard set of demands and R_i is a standard routing of the demands in D_i that uses fewer than $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers. The goal of **Block**$(C, (D_i, R_i))$ is to alter D_i and R_i to create a new standard set of demands D_{i+1} and a standard routing R_{i+1} for D_{i+1} such that the number of wavelength interchangers that service a demand with either input wavelength λ_1 or output wavelength γ_2 is strictly greater under D_{i+1} and R_{i+1} than it was under D_i and R_i. In order to achieve this, **Block**$(C, (D_i, R_i))$ will use the demands in D^H to "hold" the wavelength interchangers in W^H so that they continue to service only demands with an output fiber in $\mathcal{O}^H$. Then it will use the demands serviced by two of the wavelength interchangers in W_i^B that must be servicing two demands by Lemma 4.2 to create a new demand with input wavelength λ_1 and output wavelength γ_2 that must be serviced by a wavelength interchanger that was not in W_i^B.

Block$(C, (D_i, R_i))$

1. Take two wavelength interchangers, WI_u and WI_v, in W_i^B that, by Lemma 4.2, each service two demands.
2. **Uncross(**WI_u**)** and **Uncross(**WI_v**)**.
3. Let $(a_{u1}, \lambda_1, b_{u1}, \gamma_1)$ and $(a_{u2}, \lambda_2, b_{u2}, \gamma_2)$ be the two resulting demands that WI_u services. Let $(a_{v1}, \lambda_1, b_{v1}, \gamma_1)$ and $(a_{v2}, \lambda_2, b_{v2}, \gamma_2)$ be the two resulting demands that WI_v services.
4. Remove $(a_{u1}, \lambda_1, b_{u1}, \gamma_1)$ and $(a_{v2}, \lambda_2, b_{v2}, \gamma_2)$ from D_i and route all remaining demands according to R_i to create D_{i+1} and R_{i+1}.
5. For any demand $d_j \in D_{i+1}$ of the form $(I_e, \lambda_t, h_j, \gamma_2)$, remove d_j from D_{i+1}, replace it with the demand $d_j' = (I_e, \lambda_t, h_j, \gamma_1)$ and route d_j' along the same path that d_j was routed along.
6. Add $(a_{v2}, \lambda_2, b_{u1}, \gamma_1)$ to D_{i+1} and add a valid route for this demand to R_{i+1}.
7. For any demand $d_j' \in D_{i+1}$ of the form $(I_e, \lambda_t, h_j, \gamma_1)$, remove d_j' from D_{i+1}, replace it with demand $d_j = (I_e, \lambda_t, h_j, \gamma_2)$ and route d_j along the same path that d_j' was routed along.
8. Add $(a_{u1}, \lambda_1, b_{v2}, \gamma_2)$ to D_{i+1} and add a valid route for this demand to R_{i+1}.
9. Return (D_{i+1}, R_{i+1}).

We start with D_0 and R_0 as defined above and then inductively define $(D_{i+1}, R_{i+1}) = \mathbf{Block}(C, (D_i, R_i))$ as long as R_i uses fewer than $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers. Consider $\mathbf{Block}(C, (D_i, R_i))$. By Lemma 4.2 and the assumption that R_i uses fewer than $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers, there must exist two wavelength interchangers in W_i^B that service two demands. Given these two wavelength interchangers, we then "uncross" their demands. This does not change any of the wavelength channels used along any of the fibers in C. In particular, while this alters the set of demands, the new set of demands is still a standard set of demands. In Step 4 we remove two demands and leave the remaining demands unchanged. By our choice of demands to remove we insure that WI_v and WI_u will still each service one demand with either input wavelength λ_1 or output wavelength γ_2.

Next we change the output wavelength of every demand in D^H to be γ_1. Since R_i is a standard routing, all the output paths for demands in D^H are edge disjoint from all other demands. Therefore this step does not create an invalid set of demands or routings for those demands. Changing the output wavelength of every demand in D^H means that each wavelength interchanger in W^H is servicing a demand with output γ_1 when a route for the new demand with output wavelength γ_1 in Step 6 is found. As a result this demand can only be serviced by a wavelength interchanger in either W_i^B or W_i^F and any route for this demand will use an output path that is edge disjoint from the output paths for all demands in D^H. If the wavelength interchanger that does service the demand from Step 6 is in W_i^F, then it will be in W_{i+1}^F and clearly if it is in W_i^B then it will be in W_{i+1}^B. So after Step 6, $|W_{i+1}^B| = |W_i^B|$, $|W_{i+1}^F| = |W_i^F|$ and $|W^H|$ of course remains fixed.

Next, in Step 7, we switch all $d_j' = (I_e, \lambda_t, h_j, \gamma_1)$ back to $d_j = (I_e, \lambda_t, h_j, \gamma_2)$. Again this is possible since R_i is a standard routing and the output paths of the routes for these demands must be edge disjoint from all other output paths. Furthermore switching the output wavelength of each of the demands serviced by a wavelength interchanger in W^H back to γ_2 means that the new demand in Step 8 must be serviced by a wavelength interchanger in either W_i^B or W_i^F and the route chosen to service the demand will have an edge disjoint output path from all output paths for demands in D^H.

Since C is strictly non-blocking, a valid route for the new demand in Step 8 must exist. A wavelength interchanger in W_i^F must be used to service this demand since all wavelength interchangers in W_i^B are servicing a demand with either input wavelength λ_1 or output wavelength γ_2 and all wavelength interchangers in W^H are servicing a demand with output wavelength γ_2. Since the demand requested in Step 8 has both input wavelength λ_1 and output wavelength γ_2, the wavelength interchanger that services this new demand will be in W_{i+1}^B. Hence, after Step 8, D_{i+1} is a standard set of demands and R_{i+1} is a standard routing for those demands such that $|W^H|$ remains fixed, $|W_{i+1}^B| = |W_i^B|+1$ and $|W_{i+1}^F| = |W_i^F| - 1$. Thus we can conclude:

Lemma 4.3 *If C is a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $k_2 < n_1k_1$, D_i is a standard set of demands and R_i is a standard routing of the demands in D_i that uses fewer than $\min(k_1, nk_1 - k_2) + k_2 - 1$ wavelength interchangers, then $\mathbf{Block}(C, (D_i, R_i))$ can be executed and $(D_{i+1}, R_{i+1}) = \mathbf{Block}(C, (D_i, R_i))$ where D_{i+1} is a standard set of demands and R_{i+1} is a standard routing of the demands in D_{i+1}. Also, D_{i+1} and R_{i+1} are such that, $|W^H|$ remains fixed, $|W_{i+1}^B| = |W_i^B| + 1$ and $|W_{i+1}^F| = |W_i^F| - 1$.*

We now use Lemma 4.3 to show that any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $k_1 \leq k_2 < n_1k_1$ must have at least $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers.

Theorem 4.4 *For any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $k_1 \leq k_2 < n_1k_1$ there must be a standard set of demands D_t and a standard routing R_t of D_t that uses at least $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers.*

Proof: By contradiction. Let C be a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect where $k_1 \leq k_2 < n_1k_1$. As noted earlier, for such a cross-connect there exists a standard set of demands D_0 with standard routing R_0. Let $M = \min(k_1, n_1k_1 - k_2) + k_2 - 1$. If R_0 uses at least M wavelength interchangers then we are done. Otherwise, by Lemma 4.3, we can repeatedly apply **Block** to produce a series of pairs $(D_{i+1}, R_{i+1}) = \mathbf{Block}(C, (D_i, R_i))$ where D_i is a standard set of demands and R_i is a standard routing of the demands in D_i as long as R_i has the property that it uses fewer than M wavelength interchangers. By Lemma 4.3 we know that $|W_{i+1}^F| = |W_i^F| - 1$ because the demand in Step 9 must be routed through a wavelength interchanger in W_i^F. Thus if $|W_0^F| = t$ then either for some $i < t$, R_i uses at least M wavelength interchangers and we are done or $|W_t^F| = 0$. If R_t uses fewer than M wavelength interchangers then Lemma 4.3 guarantees that we can perform $\mathbf{Block}(C, (D_t, R_t))$. However, the demand in Step 9 must be serviced by a wavelength interchanger in W_t^F. Therefore Step 9 can not be completed in $\mathbf{Block}(C, (D_t, R_t))$. Thus it must be that R_t uses $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ or more wavelength interchangers. ■

Theorem 4.4 says that if $k_2 \leq (n_1 - 1)k_1$, then $k_1 + k_2 - 1$ wavelength interchangers are needed. However, if $(n_1-1)k_1 < k_2 < n_1k_1$ then it only says that $n_1k_1 - 1$ wavelength interchangers are needed. We now show that in this case, there must, in fact, be at least n_1k_1 wavelength interchangers.

The idea is that we would like to continue executing $(D_{i+1}, R_{i+1}) = \mathbf{Block}(C, (D_i, R_i))$ until R_{i+1} uses $n_1k_1 - 1$ wavelength interchangers and $W^F_{i+1} = 0$. However, in general, it is possible that at some point, $\mathbf{Block}(C, (D_i, R_i))$ could return a standard set of demands D_{i+1} and a standard routing R_{i+1} for D_{i+1} such that $n_1k_1 - 1$ wavelength interchangers in C are used to service the demands but $W^F_{i+1} > 0$. Once this many wavelength interchangers in C are used to service a demand we can no longer run $\mathbf{Block}(C, (D_{i+1}, R_{i+1}))$. Therefore we need a procedure that is similar in nature to $\mathbf{Block}(C, (D_{i+1}, R_{i+1}))$ but instead of using the demands corresponding to two wavelength interchangers that each service two demands it will use one such set of demands, which must exist by Lemma 4.2 and one single demand in D^H that uses input wavelength λ_1 and output wavelength γ_2. Before presenting the new construction we need to prove that such a demand in D^H exists.

Lemma 4.5 *Let C be a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect where $(n_1-1)k_1 < k_2 < n_1k_1$. If D_i is a standard set of demands, then at least one demand in D^H has input wavelength λ_1 and output wavelength γ_2.*

Proof: Since D_i is a standard set of demands, there must be one demand in D_i with output wavelength γ_2 for each output fiber in $\mathcal{O}^H$. The set of all such demands is by definition D^H. If $(n_1-1)k_1 < k_2 < n_1k_1$, then the number of output fibers in $\mathcal{O}^H$ is $n_1k_1 - z > (n_1-1)k_1$. Therefore at least one input fiber must have n_1 demands originating from it, all of which are in D^H. Since no two demands from the same input fiber can have the same input wavelength, this input fiber must have exactly one demand for each possible input wavelength. Therefore it has a demand in D^H with input wavelength λ_1. Since all demands in D^H have output wavelength γ_2, this demand is a demand in D^H with input wavelength λ_1 and output wavelength γ_2. ■

Now suppose we have a standard set of demands D_i and a standard routing R_i that routes D_i in such a way that all $n_1k_1 - 1$ wavelength interchangers in C service a demand in D_i. We can use Lemma 4.5 to show that in fact C must have n_1k_1 wavelength interchangers if it is strictly non-blocking.

First we need to define the following notation. Let p_1 and p_2 be paths such that the last node v in p_1 is the first node in p_2. We will use the notation $p_1||p_2$ to denote the path formed by following p_1 to v and then following p_2. Consider $\mathbf{Block2}(C, (D_i, R_i))$ defined as follows.

Block2$(C, (D_i, R_i))$

1. Take one wavelength interchanger, WI_v, in W^B_i that services two demands and one wavelength interchanger WI_u in W^H that services a demand in D^H with input wavelength λ_1 and output wavelength γ_2.
2. **Uncross**(WI_v).
3. Let $d_{v1} = (a_{v1}, \lambda_1, b_{v1}, \gamma_1)$ and $d_{v2} = (a_{v2}, \lambda_2, b_{v2}, \gamma_2)$ be the two resulting demands that WI_v services.
4. Let $d_u = (I_u, \lambda_1, h_2, \gamma_2)$ be the demand that WI_u services.
5. Remove d_u from D_i and replace it with demand $d'_u = (I_u, \lambda_1, h_2, \gamma_1)$ and route d'_u along the path that d_u was routed on.
6. Remove $d_{v1} = (a_{v1}, \lambda_1, b_{v1}, \gamma_1)$ and route all remaining demands according to R_i.
7. Add the demand $d'_{v1} = (a_{v1}, \lambda_1, h_2, \gamma_2)$ to D_i, add a valid route r'_{v1} for d'_{v1} to R_i and let p'_{v1} be the input path of r'_{v1}.
8. Let WI_x be the wavelength interchanger in W^F_i that services $d'_{v1} = (a_{v1}, \lambda_1, h_2, \gamma_2)$.
9. Let $d_x = (a_{x2}, \lambda_2, b_{x1}, \gamma_1)$ be the other demand that WI_x services and let p_{x1} be the output path of the route for d_x.
10. Remove d_x and d'_{v1} from D_i.
11. Add $d_q = (a_{v1}, \lambda_1, b_{x1}, \gamma_1)$ to D_i to create D_{i+1} and add the valid route $r_q = p'_{v1}||p_{x1}$ for d_x to R_i to create R_{i+1}.
12. For any demand $d_j \in D_{i+1}$ of the form $(I_e, \lambda_t, h_j, \gamma_2)$, remove d_j from D_{i+1}, replace it with the demand $d'_j = (I_e, \lambda_t, h_j, \gamma_1)$ and route d'_j along the same path that d_j was routed on.
13. Add $(a_{x2}, \lambda_2, b_{v1}, \gamma_1)$ to D_{i+1} and add a valid route for this demand to R_{i+1}.
14. For any demand $d'_j \in D_{i+1}$ of the form $(I_e, \lambda_t, h_j, \gamma_1)$, remove d'_j from D_{i+1}, replace it with demand $d_j = (I_e, \lambda_t, h_j, \gamma_2)$ and route d_j along the same path that d'_j was routed on.
15. Return (D_{i+1}, R_{i+1}).

Lemma 4.6 *If C is a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $n_1k_1 - 1$ wavelength interchangers where $(n_1-1)k_1 < k_2 < n_1k_1$, D_i is a standard set of demands and R_i is a standard routing of the demands in D_i that uses all $n_1k_1 - 1$ wavelength interchangers, then* $\mathbf{Block2}(C, (D_i, R_i))$ *can be executed and* $(D_{i+1}, R_{i+1}) = \mathbf{Block2}(C, (D_i, R_i))$ *where D_{i+1} is a standard set of demands and R_{i+1} is a standard routing of the demands in D_{i+1} such that all $n_1k_1 - 1$ wavelength interchangers in C service a demand in D_{i+1}.*

Also, D_{i+1} and R_{i+1} are such that, $|W^H|$ remains fixed, $|W^B_{i+1}| = |W^B_i| + 1$ and $|W^F_{i+1}| = |W^F_i| - 1$.

Proof: First, Lemmas 4.2 and 4.5 show the existence of a wavelength interchanger $WI_v \in W^B_i$ that services two demands and a wavelength interchanger $WI_u \in W^H$ that services a demand with input wavelength λ_1 and output wavelength γ_2. After "uncrossing" the demands that WI_v services, we change the output wavelength from γ_2 to γ_1 of the demand d_u that WI_u services. The inductive assumption that R_i is a standard routing allows us to change this output wavelength without rerouting the demand.

By the definition of a standard set of demands, no other demand in D_i used the output fiber h_2 that d_u uses. Therefore after we change the only demand in D_i that uses h_2 to a demand with output wavelength γ_1 it becomes valid to add a new demand with output wavelength γ_2 and output fiber h_2.

In Step 6 we remove the demand serviced by WI_v with input wavelength λ_1. At this point we then add a demand from input fiber a_{v1} to output fiber h_2. This new demand uses input wavelength λ_1 and output wavelength γ_2. Therefore it must be serviced by a wavelength interchanger $WI_x \in W^F_i$ since all of the wavelength interchangers in W^B_i and W^H are already servicing a demand with either input wavelength λ_1 or output wavelength γ_2. The second demand d_x that WI_x services must exist because it is assumed that every wavelength interchanger in C services a demand in D_i.

Since R_i is able to route the first half of demand d'_v along p'_{v1} with wavelength λ_1 and the second half of d_x along p_{x1} with wavelength γ_1 it must be possible to now route demand d_q along $r_q = p'_{v1}||p_{x1}$ in Step 11. Notice that this adds WI_x to W^B_{i+1}.

Next we "switch" the output wavelength used by all demands in D^H. Since R_i is assumed to be a standard routing, switching the output wavelength of the demands in D^H does not make their routes invalid. The reason for switching the output wavelength of all demands in D^H is to insure that the new demand in Step 13 is not serviced by a wavelength interchanger in W^H. Since this new demand is a valid demand and C is strictly non-blocking, there must exist a valid route for this demand. Furthermore the output path of the route for the new demand must be edge disjoint from all demands that use output wavelength γ_1. Therefore the route chosen for the new demand will maintain the inductive invariant that all demands not in D^H use edge disjoint output paths from those used by demands in D^H. After we have found a route for the demand made in Step 13 we switch the output wavelength of all demands in D^H back to γ_2 so that D_{i+1} meets the definition of a standard set of demands.

Notice that $|D_{i+1}| = k_2 + z$, $2z$ of the demands in D_{i+1} are demands with an input fiber in $\mathcal{I}^A$, an output fiber in $\mathcal{O}^A$, input wavelength either λ_1 or λ_2 and output wavelength either γ_1 or γ_2 and the other $k_2 - z$ demands in D_{i+1} have output wavelength γ_2 and an output fiber in $\mathcal{O}^H$. Furthermore the routing R_{i+1} is such that each wavelength interchanger in W^H services at most one demand and the set of output paths for demands in D^H are edge disjoint from all output paths for demands in D^A_{i+1}. Therefore R_{i+1} is a standard routing of D_{i+1}. Furthermore, $|W^H|$ remains fixed, $|W^B_{i+1}| = |W^B_i|+1$ and $|W^F_{i+1}| = |W^F_i|-1$. Finally, since **Block2**$(C, (D_i, R_i))$ only ever removed a demand from a wavelength interchanger that already serviced two demands, each wavelength interchanger in C must service a demand in D_{i+1} under R_{i+1}. ■

By Lemma 4.3 we can repeatedly perform **Block**$(C, (D_i, R_i))$ until $n_1k_1 - 1$ wavelength interchangers are used to service a standard set of demands D_i routed according to a standard routing R_i of D_i. Using Lemma 4.6 we can then repeatedly augment D_i and R_i until we arrive at a standard set of demands D_t and a standard routing R_t of D_t for which $|W^F_t| = 0$. In **Block2**$(C, (D_t, R_t))$, the new demand in Step 7 is a valid demand and therefore any strictly non-blocking WDM cross-connect must be able to find a route for this demand. However, since it must be serviced by a wavelength interchanger from W^F_t, C must not be strictly non-blocking. This contradiction leads to the following result.

Theorem 4.7 *Any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $(n_1 - 1)k_1 < k_2 < n_1k_1$, where $n_1 > 1$ and $n_2 > 1$ are the number of available wavelengths in each input fiber and output fiber respectively, will have at least n_1k_1 wavelength interchangers.*

Together the results in Theorems 4.1, 4.4 and 4.7 imply the following.

Theorem 4.8 *Any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $n_1 > 1$ input wavelengths and $n_2 > 1$ output wavelengths must have at least* $\min(k_1 + k_2 - 1, n_1k_1)$ *wavelength interchangers.*

5 Discussion

A wavelength channel is logically equivalent to a time-slot in a time division multiplexed (TDM) network [Sch87]. Thus our results for minimizing the number of wavelength interchangers in a strictly non-blocking WDM cross-connect are applicable to the problem of minimizing the number of locations in a strictly non-blocking TDM cross-connect where time-slot interchange takes place.

We are currently working on a number of variations on the problem presented here. For instance, consider the case where the number of input fibers differs from the number of output fibers but the set of input wavelengths is identical to the set of output wavelengths. In this case, more general topologies than those of WDM split cross-connects must be considered as possibilities in the design of optimal strictly non-blocking cross-connects.

A wide-sense non-blocking WDM cross-connect is one where there is a fixed routing algorithm so that any new demand can be satisfied using the same algorithm without disturbing any previously routed demands under the assumption that the previous demands were also routed using the algorithm [Spa87, Hin93]. The ideas used in the construction of a strictly non-blocking WDM cross-connect presented in Section 3 could also be used to construct $k_1 \times k_2$ wide-sense non-blocking WDM cross-connects with $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. It would be interesting to know if the number of wavelength interchangers needed for a wide-sense non-blocking WDM cross-connect is strictly less than or the same as the number needed for a strictly non-blocking WDM cross-connect.

Another area of future work is to consider the number of wavelength interchangers needed for a "generalized" notion of valid demands and routes (there may be more than one demand from a given input channel and routes for such demands can intersect over some initial path). This notion of multicasting has been considered in the case where there is only one channel available on each link and it has been shown that the upper and lower bounds on the size (i.e. the number of links in the underlying topology) of pathwise wide-sense generalized cross-connects remains $\Theta(k \log k)$ whereas the lower bound for the size of a pathwise strictly non-blocking generalized cross-connect has been shown to be $\Omega(k^2)$ [BP80].

References

[Ben35] V. E. Beněs. *Mathematical Theory of Connecting Networks and Telephone Traffic.* Academic Press, New York, NY, 1935.

[BP74] L. A. Bassalygo and M. S. Pinsker. Complexity of an optimal nonblocking switching network without reconnections. *Problems Inform. Transmission*, 9:64–66, 1974.

[BP80] L. A. Bassalygo and M. S. Pinsker. Asymptotically optimal networks for generalized rearrangeable switching and generalized switching without rearrangement. *Problemy Peredachi Informatsii*, 16:94–98, 1980.

[Can72] D. Cantor. On construction of nonblocking switching networks. In *Proc. of Symposium on Computer-Communications Networks and Teletraffic*, pages 253–255, 1972.

[Hin93] H. S. Hinton. *An Introduction to Photonic Switching Fabrics.* Plenum Press, New York, NY, 1993.

[KK99] J. Kleinberg and A. Kumar. Wavelength conversion in optical networks. In *Proceedings of SODA'99*, pages 566–575, 1999.

[Pip82] N. Pippenger. Telephone switching networks. In *Proc. AMS Symposium on Applied Mathematics*, pages 101–133, 1982.

[PV76] N. Pippenger and L. Valiant. Shifting graphs and their applications. *J. ACM*, 23:423–432, 1976.

[RS97] R. Ramaswami and G. H. Sasaki. Multiwavelength optical networks with limited wavelength conversion. In *Proceedings of IEEE INFOCOM*, volume 2, pages 489–498, 1997.

[RW00] A. Rasala and G. Wilfong. Strictly non-blocking WDM cross-connects. In *Proceedings of SODA*, pages 606–615, 2000.

[SAS98] S. Subramaniam, M. Azizoglu, and A. Somani. On the optimal placement of wavelength converters in wavelength-routed networks. In *Proc. of INFOCOM'98*, 1998.

[SB99] T. E. Stern and K. Bala. *Multiwavelength Optical Networks.* Addison-Wesley, Reading, MA, 1999.

[Sch87] M. Schwartz. *Telecommunication Networks: Protocols, Modeling, and Analysis.* Addison-Wesley, Reading, MA, 1987.

[Sha50] C. E. Shannon. Memory requirements in a telephone exchange. *Bell System Tech. J.*, 29:343–349, 1950.

[Spa87] R. A. Spanke. Architectures for guided-wave optical space switching systems. *IEEE Communications Magazine*, 25(5):42–48, 1987.

[TS99] S. Thiagarajan and A. Somani. An efficient algorithm for optimal wavelength converter placement on wavelength-routed networks with arbitrary topologies. In *Proceedings of INFOCOM'99*, 1999.

[WMDZ99] G. Wilfong, B. Mikkelsen, C. Doerr, and M. Zirngibl. WDM cross-connect architectures with reduced complexity. *Journal of Lightwave Technology*, pages 1732–1741, October 1999.

[WW98] G. Wilfong and P. Winkler. Ring routing and wavelength translation. In *Proceedings of SODA'98*, pages 333–341, 1998.