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

# 1 Summary

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 $C$. 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. A desirable property of a WDM cross-connect $C$ is for $C$ to always be capable of routing any unused wavelength channel on a fiber entering $C$ onto any unused wavelength channel on any fiber leaving $C$ regardless of the other routes used in $C$. Any such WDM cross-connect is said to be strictly non-blocking. We consider the problem of minimizing the number of wavelength interchangers needed in a strictly non-blocking WDM cross-connect. We begin by presenting a number of simple designs for WDM cross-connects with $n_1$ wavelength channels on each of $k_1$ input fibers and $n_2$ wavelength channels on each of $k_2$ output fibers and prove that they are strictly non-blocking. Together these designs provide a 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 $k_1 \leq k_2$, $n_1$ and $n_2$. Then a proof is given showing that this method is optimal in that any such WDM cross-connect must have at least $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers.

# 2 Introduction

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 so that any connections currently established within the cross-connect need not 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]. In this case, since a demand never changes wavelengths and therefore no wavelength interchangers are needed in the cross-connect, the optimization criterion

*Supported by a Lucent Technologies GRPW fellowship.

can be thought of as minimizing the number of edges in the graph representing the topology of the cross-connect. Such 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 demands.

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 possibilities 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 optimal placement of wavelength interchangers for various network topologies and traffic models to minimize the blocking probability [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 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]. The reason we consider cross-connects that have a different number of input and output fibers and a different number of input and output wavelengths is as follows. 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. 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 with $n_1$ input and $n_2$ output wavelengths that use $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. These are proved 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 $\min(n_1 k_1 - k_2, k_1) + k_2 - 1$ wavelength interchangers are required and so our construction is optimal if $k_2 \leq (n_1 - 1)k_1$. However if $n_1 k_1 > k_2 > (n_1 - 1)k_1$ 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 but since this procedure and its proof are fairly complicated we have put their details in an attached Appendix for the interested reader.

## 3 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*.

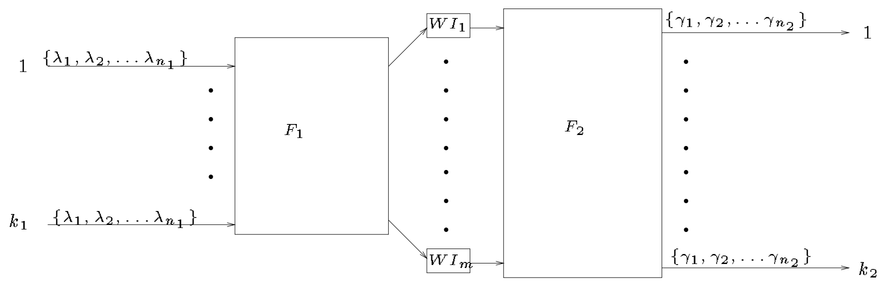


Figure 1: A split cross-cross connect with $k_1$ input fibers and $k_2$ output fibers.

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 $\lambda_i$ on a fiber directed into $WI$ and any output wavelength $\gamma_j$, $WI$ can move the signal coming in with wavelength $\lambda_i$ onto the outgoing fiber using wavelength $\gamma_j$ so long as $WI$ has not moved any other incoming signal onto wavelength $\gamma_j$. Other types of optical components in a cross-connects are input side and output side *switches*. An $n \times q$ input side switch $s$ has $n$ fibers directed into it, $q$ fibers directed out of it. For any input wavelength $\lambda_i$, a signal using wavelength $\lambda_i$ on any fiber into $s$ can be switched by $s$ onto any fiber out of $s$ using wavelength $\lambda_i$ with the restriction that at most one signal using $\lambda_i$ 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 *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, $WI_1, \ldots, WI_m$.

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

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 $R$. If $r$ is the route for $d$ used by $C$ to satisfy $d$ then we say that the fibers and devices along $r$ *service* $d$. 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. An input wavelength $\lambda_i$ (or output wavelength $\gamma_i$) is said to be *available* at a wavelength interchanger $WI_j$ if there is no demand serviced by $WI_j$ with input wavelength $\lambda_i$ (or output wavelength $\gamma_i$).

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 $\{d\} \bigcup D$ is a valid set of demands, there is a valid route $r$ of $d$ with respect to $R$.

# 4 Constructions for strictly non-blocking WDM cross-connects

We will consider two constructions of $k_1 \times k_2$ WDM cross-connects that are strictly non-blocking. The proofs that they are strictly non-blocking are omitted in this extended abstract.

As mentioned in the introduction, the problem of optimizing the design of cross-connects in the case where each input and output line carries only a single channel has been well studied [Ben35, Sha50, PV76, BP74, Pip82]. Any cross-connect of this type is said to be pathwise strictly non-blocking if there always exists a valid route from any unused input fiber to any unused output fiber that doesn't require rerouting any previously routed demands. If we replace each fiber in a traditional (i.e. non-WDM) pathwise strictly non-blocking cross-connect design with a fiber that supports $n$ wavelengths, then the result is a design for a fabric $F$ that supports $n$ wavelengths and is said to be pathwise strictly non-blocking.

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 request. Since the number of available input wavelengths is $n_1$ and the number of input fibers is $k_1$ the largest number of demands in any valid demand set is no more than $n_1 k_1$. We define $C$ to be a *dedicated cross-connect* if $C$ is a split cross-connect where $F_2$ is $n_1 k_1 \times k_2$ pathwise strictly non-blocking and the topology of $F_1$ 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.

**Theorem 4.1** *Let $C$ be a WDM dedicated cross-connect where the number of wavelength interchangers in $C$ is $n_1 k_1$. Then $C$ is strictly non-blocking.*

By Theorem 4.1, any WDM dedicated 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.

**Theorem 4.2** *Let $C$ be a $k_1 \times k_2$ WDM split cross-connect with $k_1 + k_2 - 1$ wavelength interchangers. If $F_1$ and $F_2$ are pathwise strictly non-blocking, then $C$ is strictly non-blocking.*

Theorem 4.2 shows a second construction for a strictly non-blocking $k_1 \times k_2$ WDM cross-connect. Notice that as long as $k_2 \leq (n_1 - 1)k_1$ then $k_1 + k_2 - 1 < n_1 k_1$ and therefore the second construction 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 $k_1 \times k_2$ WDM dedicated cross-connect with $n_1 k_1$ dedicated wavelength interchangers would be sufficient.

**Theorem 4.3** *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.*

# 5 Lower bounds

We begin by considering the case when $k_2 \geq n_1 k_1$ 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.

**Lemma 5.1** *If $C$ is a strictly non-blocking $k_1 \times k_2$ WDM cross-connect with $k_2 \geq n_1 k_1$, 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 above argument that any strictly non-blocking WDM cross-connect must have $k_2$ at least wavelength interchangers. In order to show that any strictly non-blocking WDM split cross-connect actually requires $\min(k_1, n_1 k_1 - k_2) + k_2 - 1$ wavelength interchangers, we will create a demand set $D$ and a routing $R$ of $D$ that routes $D$ in such a way that any strictly non-blocking split cross-connect with fewer wavelength interchangers must use each wavelength interchanger to service a demand with either input wavelength $\lambda_1$ or output wavelength $\gamma_2$. 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 $\lambda_1$ and output wavelength $\gamma_2$. 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 $\min(k_1, n_1 k_1 - k_2) + k_2 - 1$ wavelength interchangers.

Consider a strictly non-blocking $k_1 \times k_2$ WDM cross-connect $C$ where $k_2 < n_1 k_1$. Let $z = \min(k_1, n_1 k_1 - k_2)$, $A = \{a_1, a_2 \ldots a_z\}$ be a subset of $I$ consisting of $z$ input fibers, $B = \{b_1, b_2, \ldots b_z\}$ be a subset of $z$ output fibers in $O$ and $E = \{e_1, e_2 \ldots e_{k_2 - z}\}$ 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 $A$, an output fiber in $B$, input wavelength either $\lambda_1$ or $\lambda_2$ and output wavelength either $\gamma_1$ or $\gamma_2$ and

3. the other $k_2 - z$ demands in $D$ have output wavelength $\gamma_2$ and an output fiber in $E$.

For such a standard set of demands, $D$, define $D^B$ to be the subset of demands in $D$ with an output fiber in $B$ and $D^E$ to be the subset of demands in $D$ with an output fiber in $E$.

Suppose we have a routing $R$ for a standard set of demands $D$, then we can partition the set of wavelength interchangers in $C$ into three sets depending on the types of demands that each service. Let $E_{DR}$ be the set of wavelength interchangers that service a demand in $D^E$. Notice that any wavelength interchanger can service at most one demand in $D^E$ since all demands that go to an output fiber in $E$ use the same output wavelength $\gamma_2$. Therefore since $D$ is a standard set of demands, $|E_{DR}| = k_2 - z$. Let $A_{DR}$ be the set of wavelength interchangers that are not in $E_{DR}$ yet service a demand with either input wavelength $\lambda_1$ and/or output wavelength $\gamma_2$. Let $B_{DR}$ be the set of all other wavelength interchangers. If the total number of wavelength interchangers in $C$ is $m \leq k_2 + z - 2$, then $|E_{DR}| = k_2 - z$ and $|A_{DR}| + |B_{DR}| \leq 2z - 2$.

We define a valid routing $R$ of a standard set of demands $D$ to be a *standard* routing if

1. each wavelength interchanger in $E_{DR}$ services only one demand and

2. the set of output paths of the routes for demands in $D^E$ are edge disjoint from all output paths of the routes for demands in $D^B$.

Suppose $C$ is a strictly non-blocking $k_1 \times k_2$ WDM cross-connect where $k_1 \leq k_2 < n_1 k_1$. 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 $z = \min(k_1, n_1 k_1 - k_2)$ 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 $e_j \in E$ has a demand with output wavelength $\gamma_2$. By assumption, $k_2$ is small enough relative to $k_1$ and $n_1$ that there are enough available input wavelengths so that these $k_2 - z$ demands can be made without creating any unsatisfiable demands. Clearly the set of demands that use output wavelength $\gamma_2$ must all have edge disjoint output paths. Since all other demands in $D_0^B$ are routed along a path that also has a demand in $D_0^B$ with output wavelength $\gamma_2$, all demands in $D_0^B$ are routed along output paths that are edge disjoint with respect to all output paths of routes for demands in $D_0^E$. Therefore $D_0$ is a standard set of demands and the routing described is a standard routing $R_0$ for $D_0$.

**Lemma 5.2** *If $C$ is a strictly non-blocking $k_1 \times k_2$ WDM cross-connect, $D$ is a standard set of demands and $R$ is a standard routing of the demands in $D$ that uses $\min(k_1, n_1 k_1 - k_2) + k_2 - g$ wavelength interchangers, where $g > 0$, then $g$ wavelength interchangers in $A_{DR}$ will each service two demands both of whose input wavelengths are $\lambda_1$ and $\lambda_2$ and whose output wavelengths are $\gamma_1$ and $\gamma_2$.*

**Proof:** By definition, any standard routing $R$ of $D$ routes at most one demand through any wavelength interchanger in $E_{DR}$ and therefore that demand must be a demand whose output fiber is in $E$. There are $k_2 - z$ of these demands and therefore $|E_{DR}| = k_2 - z$. Since the total number of wavelength interchangers used is $z + k_2 - g$ and $|E_{DR}| = k_2 - z$, it must be the case that the number of wavelength interchangers in $A_{DR}$ and $B_{DR}$ is $2z - g$. The $2z$

demands of $D^B$ must be routed through wavelength interchangers in $A_{DR}$ or $B_{DR}$ since $R$ is a standard routing. These $2z$ demands in $D^B$ use only input wavelengths $\lambda_1$ and $\lambda_2$ and output wavelengths $\gamma_1$ and $\gamma_2$. As a result no wavelength interchanger can service more than two of these demands. This implies that $g$ wavelength interchangers in $A_{DR}$ or $B_{DR}$ must service two demands. Any such wavelength interchanger that services two of these demands must service a demand with input wavelength $\lambda_1$ and a demand with output wavelength $\gamma_2$ and therefore is by definition in $A_{DR}$. ■

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 $D$ and a standard routing $R$ of $D$ such that every wavelength interchanger in $C$ is servicing a demand with either input wavelength $\lambda_1$ or output wavelength $\gamma_2$. Notice that this is equivalent to showing that a standard set of demands $D$ and a standard routing $R$ of $D$ exists such that $E_{DR} = k_2 - z$, $A_{DR} = 2z$ and $B_{DR} = 0$.

Let $WI_i$ 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**($WI_i$) 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**($WI_i$) 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**($WI_i$) has no effect.

Next we define the operation**Construct**$(C, (D, R))$ where $C$ is a strictly non-blocking $k_1 \times k_2$ WDM cross-connect, $D$ is a standard set of demands (with respect to $C$ and some partition of its input and output fibers) and $R$ is a standard routing of the demands in $D$. The goal of **Construct**$(C, (D, R))$ is to alter $D$ and $R$ to create a new standard set of demands $D'$ and a standard routing $R'$ for $D'$ such that the number of wavelength interchangers that service a demand with either input wavelength $\lambda_1$ or output wavelength $\gamma_2$ is strictly greater under $D'$ and $R'$ than it was under $D$ and $R$. In order to achieve this **Construct**$(C, (D, R))$ will use the demands in $D^E$ to "hold" the wavelength interchangers in $E_{DR}$ so that they continue to only service demands with an output fiber in $E$. Then it will use the demands serviced by two of the wavelength interchangers in $A_{DR}$ that must be servicing two demands by Lemma 5.2 to create a new demand with input wavelength $\lambda_1$ and output wavelength $\gamma_2$ that must be serviced by a wavelength interchanger that was not in $A_{DR}$.

**Construct**$(C, (D, R))$

1. Take two wavelength interchangers, $WI_j$ and $WI_i$, in $A_{DR}$ that, by Lemma 5.2, each service two demands.
2. **Uncross(**$WI_i$) and **Uncross(**$WI_j$).
3. Let $(a_{j1}, \lambda_1, b_{j1}, \gamma_1)$ and $(a_{j2}, \lambda_2, b_{j2}, \gamma_2)$ be the two resulting demands that $WI_j$ services. Let $(a_{i1}, \lambda_1, b_{i1}, \gamma_1)$ and $(a_{i2}, \lambda_2, b_{i2}, \gamma_2)$ be the two resulting demands that $WI_i$ services.
4. Remove $(a_{j1}, \lambda_1, b_{j1}, \gamma_1)$ and $(a_{i2}, \lambda_2, b_{i2}, \gamma_2)$ from $D$ and route all remaining demands according to $R$ to create $D'$ and $R'$.
5. For any demand $d_j \in D'$ of the form $(I_i, \lambda_t, e_j, \gamma_2)$, remove $d_j$ from $D'$, replace it with the demand $d'_j = (I_i, \lambda_t, e_j, \gamma_1)$ and route $d'_j$ along the same path that $d_j$ was routed on.

6. Add $(a_{i2}, \lambda_2, b_{j1}, \gamma_1)$ to $D'$ and add a valid route for this demand to $R'$.

7. For any demand $d'_j \in D'$ of the form $(a_i, \lambda_t, e_j, \gamma_1)$, remove $d'_j$ from $D'$, replace it with demand $d_j = (a_i, \lambda_t, e_j, \gamma_2)$ and route $d_j$ along the same path that $d'_j$ was routed on.

8. Add $(a_{j1}, \lambda_1, b_{i2}, \gamma_2)$ to $D'$ and add a valid route for this demand to $R'$.

9. Return $(D', R')$.

We start with $D_0$ and $R_0$ as defined above and then inductively define $(D_{i+1}, R_{i+1}) =$**Construct**$(C, (D_i, R_i)$. Consider **Construct**$(C, (D_i, R_i))$. By Lemma 5.2 and the inductive assumption that $R_i$ uses fewer than $\min(k_1, nk_1 - k_2) + k_2 - 1$ wavelength interchangers, there must exist two wavelength interchangers in $A_{D_iR_i}$ that service two demands. Given these two wavelength interchangers we then "uncross" their demands. This does not change any of the wavelength assignments 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 $WI_i$ and $WI_j$ will still service one demand with either input wavelength $\lambda_1$ or output wavelength $\gamma_2$.

Next we change the output wavelength of every demand in $D^{E_i}$ to be $\gamma_1$. Since $R_i$ is a standard routing, all the output paths for demands in $D^{E_i}$ 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 $D^{E_i}$ means that each wavelength interchanger in $E_{D_iR_i}$ is servicing a demand with output $\gamma_1$ when a route for the new demand with output wavelength $\gamma_1$ in Step 6 is found. As a result this demand can only be serviced by a wavelength interchanger in either $A_{D_iR_i}$ or $B_{D_iR_i}$ and any route for this demand will use an output path that is edge disjoint from the output paths for all demands in $E_{D_iR_i}$. If the wavelength interchanger that does service the demand from Step 6 is in $B_{D_iR_i}$, then it will be in $B_{D_{i+1}R_{i+1}}$ and clearly if it is in $A_{D_iR_i}$ then it will be in $A_{D_{i+1}R_{i+1}}$. So after Step 6, $|A_{D_{i+1}R_{i+1}}| = |A_{D_iR_i}|$, $|B_{D_{i+1}R_{i+1}}| = |B_{D_iR_i}|$ and $|E_{D_{i+1}R_{i+1}}| = |E_{D_iR_i}|$.

Next, in Step 7, we switch all $d'_j = (I_i, \lambda_t, e_j, \gamma_1)$ back to $d_j = (I_i, \lambda_t, e_j, \gamma_2)$. 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 $E_{D_iR_i}$ back to $\gamma_2$ means that the new demand in Step 8 must be serviced by a wavelength interchanger in either $A_{D_iR_i}$ or $B_{D_iR_i}$ and the route chosen to service the demand will have an edge disjoint output path from all output paths for demands in $E_{D_iR_i}$.

Since $C$ is strictly non-blocking, a valid route for the new demand in Step 8 must exist. A wavelength interchanger in $B_{D_iR_i}$ must be used to service this demand since all wavelength interchangers in $A_{D_iR_i}$ are servicing a demand with either input wavelength $\lambda_1$ or output wavelength $\gamma_2$ and all wavelength interchangers in $E_{D_iR_i}$ are servicing a demand with output wavelength $\gamma_2$. Since the demand requested in Step 8 has both input wavelength $\lambda_1$ and output wavelength $\gamma_2$, the wavelength interchanger that services this new demand will be in $A_{D_{i+1}R_{i+1}}$. 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 $|E_{D_{i+1}R_{i+1}}| = |E_{D_iR_i}|$, $|A_{D_{i+1}R_{i+1}}| = |A_{D_iR_i}| + 1$ and $|B_{D_{i+1}R_{i+1}}| = |B_{D_iR_i}| - 1$.

**Lemma 5.3** *If $C$ is a strictly non-blocking $k_1 \times k_2$ WDM 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* **Construct**$(C, (D_i, R_i))$ *can be* *executed and* $(D_{i+1}, R_{i+1}) = \textbf{Construct}(C, (D_i, R_i))$ *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,* $|E_{D_{i+1}R_{i+1}}| = |E_{D_iR_i}|$, $|A_{D_{i+1}R_{i+1}}| = |A_{D_iR_i}| + 1$ *and* $|B_{D_{i+1}R_{i+1}}| = |B_{D_iR_i}| - 1$.

Lemma 5.3 says that **Construct** can be repeatedly applied until we eventually arrive at a contradiction when no available wavelength interchangers remain to service the valid demand created in Step 8 of $\textbf{Construct}(C(D_i, R_i))$.

**Theorem 5.4** *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* $D$ *and a standard routing* $R$ *of* $D$ *that uses at least* $\min(n_1k_1 - k_2, k_1) + k_2 - 1$ *wavelength interchangers.*

Theorem 5.4 says that if $k_2 \leq (n_1 - 1)k_1$ then $k_1 + k_2 - 1$ wavelength interchangers are needed. However if $n_1k_1 > k_2 > (n_1 - 1)k_1$ then it only says that $n_1k_1 - 1$ wavelength interchangers are needed. In an appendix we show that in this case, then, in fact, there must be at least $n_1k_1$ wavelength interchangers.

**Theorem 5.5** *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* $n_2$ *are the number of available wavelengths in each input fiber and output fiber respectively, will have at least* $n_1k_1$ *wavelength interchangers.*

# 6 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.

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 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 4 could also be used to construct $k_1 \times k_2$ wide-sense non-blocking WDM cross-connects with $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ 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 and routes for such demands can intersect over some initial path). This problem has been considered in the case where there is only one channel available in each link and it has been shown that the upper and lower bounds on the size 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.

[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*, 2000.

[SAS98] S. Subramaniam, M. Azizoglu, and A. Somani. On the optimal placement of wavelength converters in wavelength-routed networks. In *Proceedings 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*, October 1999.

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

# A Appendix

Now we will show that any $k_1 \times k_2$ strictly non-blocking WDM cross-connect with $k_2 \leq (n_1 - 1)k_1$ must have at least $n_1 k_1$ wavelength interchangers. Note that by Lemma 5.4, if $C$ has exactly $n_1 k_1 - 1$ wavelength interchangers, then a standard set of demands $D$ and a standard routing $R$ of $D$ exists such that all $n_1 k_1 - 1$ wavelength interchangers are servicing demands in $D$.

The idea is that we would like to continue executing $(D_{i+1}, R_{i+1}) =$**Construct**$(C, (D_i, R_i))$ until $R_{i+1}$ uses $n_1 k_1 - 1$ wavelength interchangers and $B_{D_{i+1}, R_{i+1}} = 0$. However, in general, it is possible that at some point, **Construct**$(C, (D_i, R_i))$ could return a standard set of demands $D_{i+1}$ and a standard routing $R_{i+1}$ for $D_{i+1}$ such that $n_1 k_1 - 1$ wavelength interchangers in $C$ are used to service the demands but $B_{D_{i+1} R_{i+1}} > 0$. Once this many wavelength interchangers in $C$ are used to service a demand we can no longer run **Construct**$(C, (D_{i+1}, R_{i+1}))$. Therefore we need a procedure that is similar in nature to **Construct**$(C, (D, R))$ 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 5.2 and one single demand in $D^E$ that uses input wavelength $\lambda_1$ and output wavelength $\gamma_2$. Before presenting the new construction we need to prove that such a demand in $D^E$ exists.

**Lemma A.1** *Let $C$ be a strictly non-blocking $k_1 \times k_2$ WDM cross-connect where $(n_1 - 1)k_1 < k_2 < n_1 k_1$. If $D$ is a standard set of demands, then at least one demand in $D^E$ has input wavelength $\lambda_1$ and output wavelength $\gamma_2$.*

**Proof:** Since $D$ is a standard set of demands, there must be one demand in $D$ with output wavelength $\gamma_2$ for each output fiber in $E$. The set of all such demands is by definition $D^E$. If $(n_1 - 1)k_1 < k_2 < n_1 k_1$ then the number of output fibers in $B$ is $z = n_1 k_1 - k_2 < k_1$ and the number of output fibers in $E$ is $n_1 k_1 - z > (n_1 - 1)k_1$. Therefore at least one input fiber must have $n_1$ demands originating from it, all of which are in $D^E$. 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 $D^E$ with input wavelength $\lambda_1$. Since all demands in $D^E$ have output wavelength $\gamma_2$, this demand is a demand in $D^E$ with input wavelength $\lambda_1$ and output wavelength $\gamma_2$. ■

If $p_1$ and $p_2$ are paths such that the last node $v$ in $p_1$ is the first node in $p_2$ then in what follows, we will use the notation $p_1 || p_2$ to denote the path formed by following $p_1$ to $v$ and then following $p_2$.

Suppose we have a standard set of demands $D$ and a standard routing $R$ that routes $D$ in such a way that every wavelength interchanger in $C$ services a demand in $D$. Consider **Construct2**$(C, (D, R))$ defined as follows.

**Construct2**$(C, (D, R))$

1. Take one wavelength interchanger, $WI_i$, in $A_{DR}$ that services two demands and one wavelength interchanger $WI_j$ in $E_{DR}$ that services a demand in $D^E$ with input wavelength $\lambda_1$ and output wavelength $\gamma_2$.

2. **Uncross**($WI_i$).

3. Let $d_{i1} = (a_{i1}, \lambda_1, b_{i1}, \gamma_1)$ and $d_{i2} = (a_{i2}, \lambda_2, b_{i2}, \gamma_2)$ be the two resulting demands that $WI_i$ services.

4. Let $d_1 = (I_1, \lambda_1, e_2, \gamma_2)$ be the demand that $WI_j$ services.
5. Remove $d_1$ from $D$ and replace it with demand $d_1' = (I_1, \lambda_1, e_1, \gamma_1)$ and route $d_1'$ along the path that $d_1$ was routed on.
6. Remove $d_{i1} = (a_{i1}, \lambda_1, b_{i1}, \gamma_1)$ and route all remaining demands according to $R$.
7. Add the demand $d_{i1}' = (a_{i1}, \lambda_1, e_2, \gamma_2)$ to $D$ to create $D'$, add a valid route $r_{i1}'$ for $d_{i1}'$ to $R$ to create $R'$ and let $p_{i1}'$ be the input path of $r_{i1}'$.
8. Let $WI_k$ be the wavelength interchanger in $B_{D'R'}$ that services $d_{i1}' = (a_{i1}, \lambda_1, e_2, \gamma_2)$.
9. Let $d_k = (a_{k2}, \lambda_2, b_{k1}, \gamma_1)$ be the other demand that $WI_k$ services and let $p_{k1}$ be the output path of the route for $d_k$.
10. Remove $d_k$ and $d_{i1}'$ from $D'$.
11. Add $d_q = (a_{i1}, \lambda_1, b_{k1}, \gamma_1)$ to $D$ and add the valid route $r_q = p_{i1}'|WI_k|p_{k1}$ for $d_k$ to $R'$.
12. For any demand $d_j \in D'$ of the form $(I_i, \lambda_t, e_j, \gamma_2)$, remove $d_j$ from $D'$, replace it with the demand $d_j' = (I_i, \lambda_t, e_j, \gamma_1)$ and route $d_j'$ along the same path that $d_j$ was routed on.
13. Add $(a_{k2}, \lambda_2, b_{i1}, \gamma_1)$ to $D'$ and add a valid route for this demand to $R'$.
14. For any demand $d_j' \in D'$ of the form $(a_i, \lambda_t, e_j, \gamma_1)$, remove $d_j'$ from $D'$, replace it with demand $d_j = (a_i, \lambda_t, e_j, \gamma_2)$ and route $d_j$ along the same path that $d_j'$ was routed on.
15. Return $(D', R')$.

**Lemma A.2** *If $C$ is a strictly non-blocking $k_1 \times k_2$ WDM 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* **Construct2**$(C, (D_i, R_i))$ *can be executed and* $(D_{i+1}, R_{i+1}) =$ **Construct2**$(C, (D_i, R_i))$ *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, $|E_{D_{i+1}R_{i+1}}| = |E_{D_iR_i}|$, $|A_{D_{i+1}R_{i+1}}| = |A_{D_iR_i}| + 1$ and $|B_{D_{i+1}R_{i+1}}| = |B_{D_iR_i}| - 1$.*

**Proof:** First, Lemmas 5.2 and A.1 show the existence of a wavelength interchanger $WI_i \in A_{D_iR_i}$ that services two demands and a wavelength interchanger $WI_j \in E_{D_iR_i}$ that services a demand with input wavelength $\lambda_1$ and output wavelength $\gamma_2$. After "uncrossing" the demands that $WI_i$ services, we change the output wavelength from $\gamma_2$ to $\gamma_1$ of the demand $d_1$ that $WI_j$ 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 $e_2$ that $d_1$ uses. Therefore after we change the only demand in $D_i$ that uses $e_2$ to a demand with output wavelength $\gamma_1$ it becomes valid to add a new demand with output wavelength $\gamma_2$ and output fiber $e_2$.

In Step 6 we remove the demand serviced by $WI_i$ with input wavelength $\lambda_1$. At this point we then add a demand from input fiber $a_{i1}$ to output fiber $e_2$. This new demand uses input wavelength $\lambda_1$ and output wavelength $\gamma_2$. Therefore it must be serviced by a wavelength interchanger $WI_k \in B_{D_iR_i}$ since all of the wavelength interchangers in $A_{D_iR_i}$ and $E_{D_iR_i}$ are already servicing a demand with either input wavelength $\lambda_1$ or output wavelength $\gamma_2$. The second demand $d_k$ that $WI_k$ services must exist because it is assumed that every wavelength interchanger in $C$ services a demand in $D_i$.

Since $R'$ is able to route the first half of demand $d_1$ along $r_{i1}$ with wavelength $\lambda_1$ and the second half of $d_k$ along $r_{k1}$ with wavelength $\gamma_1$ it must be possible to now route demand $d_q$ along $r_q = r_{i1}||r_{k1}$ in Step 11. Notice that this adds $WI_k$ to $A_{D_{i+1}R_{i+1}}$.

Next we "switch" the output wavelength used by all demands in $D^{E_i}$. Since $R_i$ is assumed to be a standard routing, switching the output wavelength of the demands in $D^{E_i}$ does not make their routes invalid. The reason for switching the output wavelength of all demands in $D^{E_i}$ is to insure that the new demand in Step 13 is not serviced by a wavelength interchanger in $E_{D_iR_i}$. 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 $\gamma_1$. Therefore the route chosen for the new demand will maintain the inductive invariant that all demands not in $D^{E_{i+1}}$ use edge disjoint output paths from those used by demands in $D^{E_{i+1}}$. After we have found a route for the demand made in Step 13 we switch the output wavelength of all demands in $E^{D_{i+1}}$ back to $\gamma_2$ 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 $A$, an output fiber in $B$, input wavelength either $\lambda_1$ or $\lambda_2$ and output wavelength either $\gamma_1$ or $\gamma_2$ and the other $k_2 - z$ demands in $D_{i+1}$ have output wavelength $\gamma_2$ and an output fiber in $E$. Furthermore the routing $R_{i+1}$ is such that each wavelength interchanger in $E_{D_{i+1}R_{i+1}}$ services at most one demand and the set of output paths for demands in $D^{E_{i+1}}$ are edge disjoint from all output paths for demands in $D^{B_{i+1}}$. Therefore $R_{i+1}$ is a standard routing of $D_{i+1}$. Furthermore, $|E_{D_{i+1}R_{i+1}}| = |E_{D_iR_i}|$, $|A_{D_{i+1}R_{i+1}}| = |A_{D_iR_i}| + 1$ and $|B_{D_{i+1}R_{i+1}}| = |B_{D_iR_i}| - 1$. Finally, since **Construct2**$(C, (D, R))$ 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 5.3 we can repeatedly perform **Construct**$(C, (D, R))$ until $n_1k_1 - 1$ wavelength interchangers are used to service a standard set of demands $D$ and a standard routing $R$ of $D$. Using Lemma A.2 we can then repeatedly augment $D$ and $R$ until we arrive at a standard set of demands $D_i$ and a standard routing $R_i$ of $D_i$ for which $|B_{D_iR_i}| = 0$. Consider **Construct2**$(C, (D_i, R_i))$. 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 $B_{D_iR_i}$, $C$ must not be strictly non-blocking. This contradiction leads to Theorem 5.5 presented in Section 5.