Strictly Non-blocking WDM Cross-connects

April Rasala* Gordon Wilfong†

Abstract

A crucial component in wavelength division multiplexed (WDM) optical networks will be WDM cross-connects that permit a signal on any input wavelength channel to be routed onto any output wavelength channel. Such cross-connects will require wavelength interchangers, devices that have the ability to change the wavelength of a signal along its route. Due to the expected high cost of wavelength interchangers, an important design goal for WDM cross-connects is to use as few wavelength interchangers as possible. We describe a family of $k \times k$ WDM cross-connects that are strictly non-blocking in terms of both paths and wavelengths yet require only $2k - 1$ wavelength interchangers and prove that this is optimal for such strictly non-blocking WDM cross-connects.

1 Introduction

A wavelength division multiplexed (WDM) cross-connect is a directed network of fibers connected to various optical components that allow for connecting a set of *input fibers* to a set of *output fibers*. Each fiber in the network can support some fixed number, say n, of wavelength channels. That is, at any time there can be up to n signals along a fiber each using a distinct wavelength. We consider only *wavelength interchanging WDM cross-connects* meaning that we allow for the connection of a wavelength channel on an input fiber to a (possibly) different wavelength channel on an output fiber. Of course, this implies that within the cross-connect there must be devices that can switch an incoming wavelength channel onto any (possibly) different wavelength channel on an outgoing fiber. Such devices are called *wavelength interchangers* [SB99]. The other type of component found in a cross-connect is called an *optical switch* or sometimes a *wavelength selective cross-connect*. An optical switch has an arbitrary number of fibers into and out of it and any wavelength channel on any incoming fiber can be switched to the same wavelength channel on any outgoing fiber (assuming the wavelength channel is not already being used). The use of optical switches with two incoming and two outgoing fibers in the design of WDM cross-connects has been studied [DMR+99].

The problem of satisfying a request for a connection in a WDM cross-connect has two aspects to it. First a route must be found in the cross-connect from the requested input fiber to the requested output fiber. Secondly, for each fiber in the route, an unused wavelength channel must be assigned so that (1) on the input and output fibers the wavelength channels assigned are the requested ones and (2) the wavelength channels assigned on any two consecutive fibers in the route must be the same unless there is a wavelength interchanger connecting the fibers.

A number of cross-connects have recently been described and their non-blocking properties analyzed [WMDZ99]. Most of those described in [WMDZ99] are *rearrangeably non-blocking* meaning that requests for new connections may require changing the paths and/or the wavelength channels of already configured connections. In a WDM cross-connect, disrupting connections in order to create new connections is undesirable since this requires buffering the connections being rearranged, creating a costly problem [SB99]. One of the designs in [WMDZ99] has the more desirable property of being *strictly non-blocking* (meaning that it can always handle new requests for connections without disturbing those already in place) but it requires $k \log k$ wavelength interchangers for a cross-connect with k input and output fibers. Since the dominating cost of a cross-connect is generally agreed to be the cost of the wavelength interchangers, our goal will be to study the problem of minimizing the number of wavelength interchangers needed for a strictly non-blocking cross-connect. In particular, we show that such $k \times k$ cross-connects can be designed using only $2k-1$ wavelength interchangers and that this is optimal.

2 Non-blocking properties

In this section we review the definitions of various non-blocking properties for cross-connects. We begin by considering the standard definitions given for the case where there is only one (wavelength) channel available. These are sometimes called *space domain cross-connects*. There is a vast literature on the problem of space domain cross-connect design. For an overview

*MIT Laboratory for Computer Science, Cambridge, MA 02139. This work was done during a summer internship at Bell Labs/Lucent Technologies. Also supported by a Lucent GRPW Fellowship. Email: arasala@theory.lcs.mit.edu

†Bell Labs, Lucent Technologies, Murray Hill NJ 07974. Email: gtw@research.bell-labs.com.

of various space domain cross-connect architectures and their non-blocking properties see [Hui90]. In the space domain, a request for a connection requires a route in the cross-connect from the specified input fiber to the specified output fiber such that the route is edge disjoint from all currently routed connections. It is assumed that at any one time there is at most one connection request originating at any given input fiber and at most one terminating at any given output fiber. Space domain cross-connects are classified according to the strength of their non-blocking properties. A space domain cross-connect can be

(i) *pathwise rearrangeably non-blocking*, meaning that any set of connection requests can be routed through the cross-connect but any additional requests received after routing the original set of requests may require some of the previously routed requests to be re-routed,

(ii) *pathwise wide-sense non-blocking*, meaning that there is a routing algorithm such that for any sequence of connection requests and withdrawals, the connection requests can be routed using the algorithm without disturbing any of the currently routed requests or

(iii) *pathwise strictly non-blocking*, meaning that any set of requests can be routed through the cross-connect and any additional requests can be routed without disturbing the routes of the others no matter how the routes were chosen.

In the WDM setting, a request for a connection requires not only a route from the input fiber to the output fiber but also a wavelength channel assignment along the route that only changes wavelength channels at wavelength interchangers and begins and ends on the requested wavelength channels. We call these requests for connections between wavelength channels on input and output fibers, *demands*. In this case, there are two types of non-blocking characteristics to study, the pathwise non-blocking characteristics and the wavelength non-blocking characteristics. That is, when an additional demand is given, the previously routed demands may require their routes to be changed, their wavelength channel assignments to be changed or both. The definitions of *wavelength rearrangeably non-blocking*, *wavelength wide-sense non-blocking* and *wavelength strictly non-blocking* are analogous to those for pathwise non-blocking given above.

Of course, the most desirable WDM cross-connect would be one that is both pathwise and wavelength strictly non-blocking. One that is wide-sense non-blocking would also be useful assuming that the algorithm to do the routing and wavelength assignment was simple and fast. Throughout the rest of the paper we say that a WDM cross-connect is *strictly non-blocking* if it is both pathwise and wavelength strictly non-blocking. The definitions of a *wide-sense non-blocking* and *rearrangeably non-blocking* WDM cross-connect are analogous.

3 Definitions

More formally, we define a $k \times k$ WDM cross-connect supporting $n > 1$ wavelengths to be a directed acyclic graph $C = (V, A, \Lambda)$ where V is the set of nodes, A the set of arcs between the nodes and $\Lambda = \{\lambda_1, \lambda_2, \ldots, \lambda_n\}$ is the set of available wavelengths. We will usually refer to an arc as a *fiber* to be consistent with the literature in the optical networking community but it should be remembered that each fiber has a single direction along which signals are permitted to flow. The node set V is partitioned into four subsets; I the set of *input nodes*, O the set of *output nodes*, S the set of *switches* and W the set of *wavelength interchangers*. Sets I and O each contain k nodes. Each node in I has indegree 0 and outdegree 1 whereas each node in O has outdegree 0 and indegree 1. The arc directed out of a node in I is called an *input fiber* and the arc directed into a node in O is called an *output fiber*. A node in W has indegree 1 and outdegree 1 while the indegree and outdegree of a node in S are unconstrained although in current practice they are likely to have input degree and output degree equal to two. The topology of a cross-connect as given by the underlying directed graph is sometimes called its *fabric*.

A *demand* d is defined as a 4-tuple (w, x, y, z) where w is an input node, x is a wavelength, y is an output node and z is a wavelength. The wavelengths x and z will be referred to as the *input* and *output wavelengths*, respectively. A *route* r in C is a directed path from a node in I to a node in O. Along each of its fibers, r is assigned one of the n wavelengths such that consecutive fibers are assigned the same wavelength unless the common node of the fibers is in W. We sometimes say that a route is from an input fiber to an output fiber rather than from the corresponding input node to the corresponding output node. A route for a demand $d = (w, x, y, z)$ then is a route from input node w to output node y such that on the corresponding input fiber the route is assigned wavelength x and on the corresponding output fiber the route is assigned wavelength z.

A *valid demand set* is a set of demands that satisfies the following conditions:

(i) for each input node a and each wavelength λ, there is at most one demand with both a as the input node and λ as the input wavelength and

(*ii*) for each output node b and each wavelength λ, there is at most one demand with both b as the output node and λ as the output wavelength.

A demand set $D = \{d_1, d_2, \ldots, d_m\}$ is said to be *satisfied* by a cross-connect C if there exists a set of routes $R = \{r_1, r_2, \ldots, r_m\}$ where

(*i*) r_i is a route for d_i, $1 \leq i \leq m$, and

(*ii*) if for some $i \neq j$, r_i and r_j share some fiber f then they must be assigned distinct wavelengths along f.

We refer to such a set R of routes as a *valid routing* of the demand set D and we say that R *satisfies* D. If demand $d \notin D$ is such that $D \bigcup \{d\}$ is a valid demand set and R is a valid routing for D, then we say that r is a *valid route* for d (with respect to R) if $R \bigcup \{r\}$ is a valid routing for $D \bigcup \{d\}$.

We say that a wavelength interchanger WI_i *services a demand* d if d is routed through WI_i.

4 Previously studied architectures

As mentioned, the non-blocking properties of a number of WDM cross-connects were discussed in [WMDZ99]. All but one of these cross-connects were either wavelength or pathwise rearrangeably non-blocking. Many of the designs discussed were from the family of WDM cross-connects that have the general form as shown in Figure 1(a) and were called *standard design cross-connects*. The part of the fabric between the input and output fibers, labeled by F, of such a cross-connect was assumed to contain no devices for wavelength interchanging. These cross-connects were defined in the context of demands that did not specify the wavelength on the output fiber. If we add a wavelength interchanger to each of the output fibers as shown in Figure 1(b), then they can handle demands as defined in this paper where each demand specifies a particular output wavelength and the resulting cross-connect will have the same non-blocking characteristics as described in [WMDZ99]. We will refer to such cross-connects as *modified standard design cross-connects*. We assume in the following that the fabric F of a modified standard design cross-connect is at least pathwise rearrangeably non-blocking.

The problem with a modified standard design cross-connect is that, as we show below, at best it can be wavelength rearrangeably non-blocking. The idea of the argument is to first show that the existence of a wavelength wide-sense (or strictly) non-blocking modified standard design cross-connect implies the existence

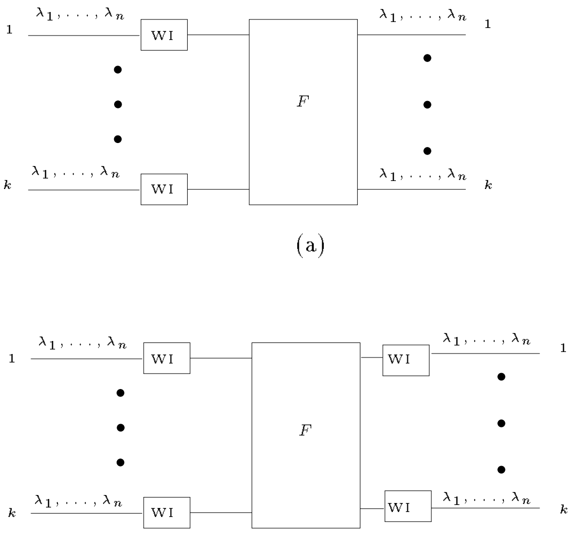

(a)

(b)

Figure 1: Standard and modified standard design cross-connects.

of an on-line algorithm for optimally edge coloring a class of bipartite multigraphs. We then show that there can be no such on-line algorithm and hence there can be no wavelength wide-sense (or strictly) non-blocking modified standard design cross-connect.

Consider the problem of optimally edge coloring a bipartite multigraph [Ber73]. That is, given a bipartite multigraph, the goal is to color the edges so that no two edges with a common endpoint share the same color. It is well known that if n is the maximum degree of any node in a bipartite multigraph $B = (B_1, B_2, E)$, then the edges in E can be optimally colored using n colors [Ber73, Gab76] and in fact a number of polynomial time algorithms have been given for the off-line version of this problem [Gab76, GK82, CH82]. An algorithm for the on-line version of this problem would work as follows. The algorithm is given the nodes of the multigraph and n, the maximum degree of any node, but the edges are only presented one at a time. Thus the initial input to the algorithm is just a description of a whole class of bipartite multigraphs; namely, those with the given numbers of nodes in its two node sets and having the given maximum degree. Then the on-line algorithm must assign a color from one of the n colors to the latest edge e presented so that the color assigned has not already been assigned to any edge incident to either endpoint of e. This must be done without changing the color of any of the previously colored edges and of course, without any knowledge of what the remaining edges will be.

Let $MG(k,n)$ be the class of bipartite multigraphs of the form $B = (B_1, B_2, E)$ where $|B_1| = |B_2| = k > 2$ and the maximum degree of any node is $n > 1$. We now show that there must be an on-line edge coloring algorithm for $MG(k,n)$ if there is a wavelength wide-sense (or strictly) non-blocking $k \times k$ modified standard design cross-connect with n wavelengths.

LEMMA 4.1. *If there is a wavelength wide-sense (or strictly) non-blocking $k \times k$ modified standard design cross-connect with $n > 1$ wavelengths and $k > 2$ then there is an on-line algorithm for optimally edge coloring the class $MG(k,n)$.*

Proof. Let C be a $k \times k$ modified standard design cross-connect with n available wavelengths and suppose C is wavelength wide-sense non-blocking. We show how using the algorithm for assigning wavelengths in C gives us an on-line algorithm for optimally edge coloring $MG(k,n)$. Let $B = (B_1, B_2, E) \in MG(k,n)$. We label the input fibers of C and the nodes of B_1 as $b_1^1, \ldots, b_k^1$. Similarly, label the output fibers of C and the nodes of B_2 as $b_1^2, \ldots, b_k^2$. Let $e_1, e_2, \ldots, e_m$ be any sequence of edges in E where $e_i = \{b_i^1, b_i^2\}$. Then demands $d_1, \ldots, d_m$ are presented to C in order one at a time where the input and output fibers of d_i are b_i^1 and b_i^2 respectively. Note that input and output wavelengths of the demands d_i can be chosen arbitrarily since there will be no more than n demands having any given input or output fiber. Since C is wavelength wide-sense non-blocking (and at least pathwise rearrangeably non-blocking), there must be a valid routing of these demands. Note that a valid route for each such demand requires assigning a constant wavelength on the part of the path in fabric F. Clearly the on-line algorithm C used to choose the wavelength assigned to demand d_i within the fabric F of C can be used to color the edges e_i of E in the same on-line fashion since it uses at most n wavelengths (colors) and cannot assign the same wavelength to more than one demand routed through a common wavelength interchanger. ■

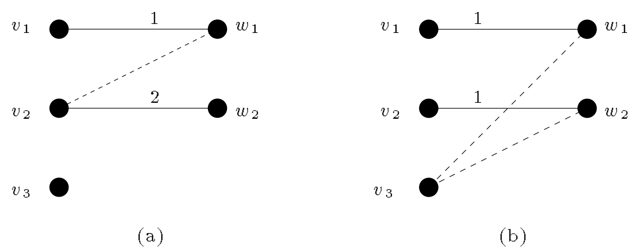

Figure 2: Failure of any on-line edge coloring algorithm.

LEMMA 4.2. *There is no on-line optimal edge coloring algorithm for $MG(k,n)$ if $k > 2$ and $n > 1$.*

Proof. Suppose such an algorithm A exists. Let $B = (B_1, B_2, E)$ be any bipartite multigraph in $MG(k,n)$ and label the nodes of B so that $\{v_1, v_2, v_3\} \subseteq B_1$, $\{w_1, w_2\} \subset B_2$. Furthermore, suppose E is such that $E_1 = \{e_{1,1}, e_{1,2}, \ldots, e_{1,n-1}\} \subset E$ is a multi-set of $n-1$ edges between v_1 and w_1. As each of these $n-1$ edges $e_{1,i}$ is presented to A, A is forced to assign a new color to $e_{1,i}$. Therefore after all $n-1$ edges between v_1 and w_1 have been presented to A, there can be at most one of the n colors not assigned to any edge adjacent to w_1. Likewise suppose E is such that $E_2 = \{e_{2,1}, e_{2,2}, \ldots, e_{2,n-1}\} \subset E$ is a multi-set of $n-1$ edges between v_2 and w_2. After all of these $n-1$ edges have been presented there will be at most one color not assigned to any of these edges and therefore only one color available for any new edges presented that are adjacent to either v_2 or w_2. In Figure 2 we show the multi-set E_1 as one edge $\{v_1, w_1\}$ and label it with the one color that is not assigned to any of the edges in E_1. Similarly the multi-set of edges E_2 is illustrated as the edge $\{v_1, w_2\}$ and labeled with the one color not assigned to any edge in E_2.

Suppose that the only color not assigned to an edge in E_1 is color 1 and the only color not assigned to an edge in E_2 is a different color, say color 2. Then as shown in Figure 2(a), A will fail if E contains the edge $e_3 = \{v_2, w_1\}$ since when e_3 is presented to A the colors already assigned to the edges in E_1 will require it to have color 1 while the colors assigned to the edges in E_2 will require it to have color 2.

On the other hand suppose that the one color not assigned to any edge in E_1, say color 1, is the same as the one color not assigned to any edge in E_2. Suppose now that E contains the edges $e_4 = \{v_3, w_1\}$ and $e_5 = \{v_3, w_2\}$ as shown in Figure 2(b) rather than the edge e_3. Since w_1 has only color 1 available, e_4 must be assigned color 1. Similarly, w_2 having only color 1 available forces e_5 to be assigned color 1. However, assigning e_4 and e_5 to both have color 1 means that v_3 has two edges adjacent to it that use the same color. Thus A fails in either case proving that no on-line algorithm can exist for optimally edge coloring this class of bipartite multigraphs. ■

We now have the results necessary to show that a modified standard cross-connect can not be better than wavelength rearrangeably non-blocking.

THEOREM 4.3. *For any $k > 2$ and $n > 1$, there is no $k \times k$ wavelength wide-sense (or strictly) non-blocking modified standard cross-connect C with n available wavelengths.*

Proof. By Lemma 4.1, if there was such a C then there would be an on-line optimal edge coloring algorithm for the class of bipartite multigraphs $MG(k,n)$. But by Lemma 4.2, there can be no such on-line edge coloring algorithm for $MG(k,n)$ where $k > 2$ and $n > 1$. ∎

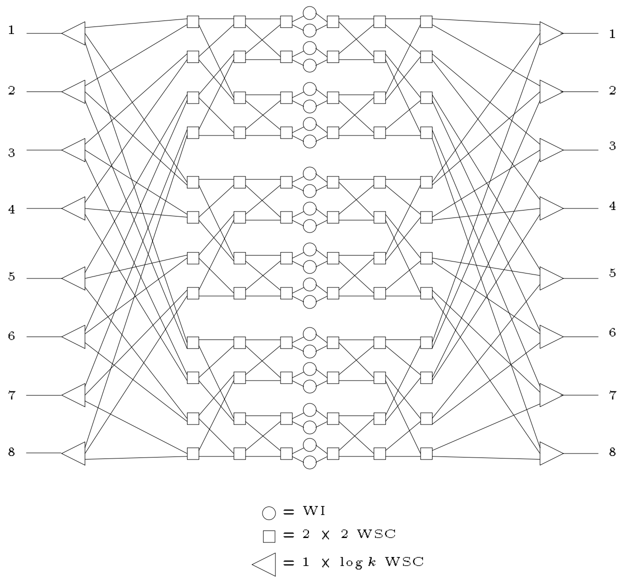

Figure 3: Cantor/2-WI-Cantor/2 cross-connect for $k = 8$.

A technique used to remove this limitation on the strength of the wavelength non-blocking characteristics of modified standard design cross-connects is to move the wavelength interchangers from the edges of the cross-connect into the "middle" [WMDZ99]. This resulted in cross-connects that were wavelength strictly non-blocking and in fact the design called Cantor/2-WI-Cantor/2 (see Figure 3) was also shown to be pathwise strictly non-blocking. Unfortunately this versatility was at the cost of using $k \log k$ wavelength interchangers rather than the $2k$ used in modified standard designs. Our goal in the next section is to show that the strictly non-blocking property can be achieved using far fewer wavelength interchangers.

5 A family of strictly non-blocking cross-connects

We introduce a new family of WDM cross-connects that have the basic form as shown in Figure 4 where the fabric of the cross-connect is split in two pieces that are separated by a level of wavelength interchangers. WDM cross-connects with such a form will be called *split cross-connects.* The only restriction on the two pieces of the fabric F_1 and F_2 is that they cannot contain any device to change the wavelength of any signal. That is, the fabric of a split cross-connect is split into two *wavelength selective* sections F_1 and F_2 connected to each other via wavelength interchangers. In a split cross connect, any directed path from an input fiber to an output fiber will pass through exactly one wavelength interchanger. Thus the only place that a route can change wavelengths is at the one wavelength interchanger that it passes through in the middle of the cross-connect. Therefore a route for a demand $(a, \lambda_1, b, \lambda_2)$ will be assigned wavelength λ_1 from the input fiber until it reaches a wavelength interchanger and from that point until the output fiber it will be assigned λ_2. Thus the wavelength assignment for any demand in such a cross-connect is completely determined by the demand. So if a split cross-connect C can satisfy any demand set then it must be that C is in a trivial sense wavelength strictly non-blocking.

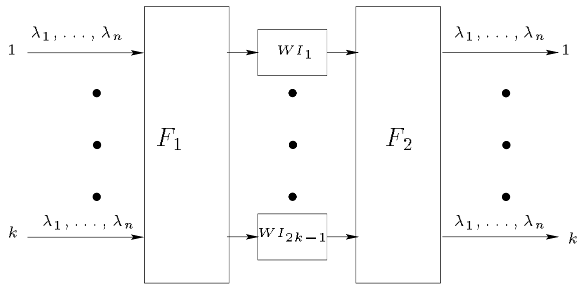

Figure 4: Split cross-connect design.

THEOREM 5.1. *Let C be a split cross-connect where the number of input and output fibers is k and the number of wavelength interchangers is $2k-1$. Suppose the topology of F_1 is that of some $k \times (2k-1)$ space domain cross-connect that is pathwise strictly non-blocking. Similarly F_2 has the topology of a $(2k-1) \times k$ space domain cross-connect that is pathwise strictly non-blocking. Then C is pathwise* and *wavelength strictly non-blocking.*

Proof. To show that C is strictly non-blocking if F_1 and F_2 are pathwise strictly non-blocking, all we need to show is that for any set D of previously routed demands and any new demand $d = (a, \lambda_1, b, \lambda_2)$ where $D \bigcup \{d\}$ is a valid demand set, there is a valid route for d. Let R be any valid routing of D.

Let D_1 be the subset of demands in D that have input wavelength λ_1 and define $W_1 \subset \{WI_1, WI_2, \ldots, WI_{2k-1}\}$ to be such that $WI_j \in W_1$ if and only if there is a demand in D_1 routed through WI_j by R. In any valid demand set there can be at most k demands that use input wavelength λ_1. Therefore the number of demands in D_1 is at most $k-1$ and $|W_1| \leq k-1$. Let W_2 denote the set of all wavelength interchangers that service a demand with output wavelength λ_2. By the same argument $|W_2| \leq k-1$.

Since there are $2k-1$ wavelength interchangers and since $|W_1|+|W_2| \leq 2k-2$ there must be some $WI_j \notin W_1 \bigcup W_2$. Since F_1 and F_2 are pathwise strictly non-blocking there must be a path from input fiber a to WI_j and a path from WI_j to output fiber b. Furthermore a path from input fiber a to WI_j can be chosen that is edge disjoint from all other paths that service a demand with input wavelength λ_1. Likewise for the path from WI_j to output fiber b. Therefore d can use this path with wavelength λ_1 from a to WI_j and wavelength λ_2 from WI_j to b without requiring that any routes in R be changed. This implies that C is a strictly non-blocking cross-connect if F_1 and F_2 are both pathwise strictly non-blocking. ∎

Notice that an analogous argument to the one given in the proof of Theorem 5.1 could also be used to show that C is pathwise wide-sense non-blocking and wavelength strictly non-blocking if F_1 and F_2 are pathwise wide-sense non-blocking.

THEOREM 5.2. *Let C be a split cross-connect where the number of input and output fibers is k and the number of wavelength interchangers is $2k-1$. Suppose the topology of F_1 is that of some $k \times (2k-1)$ space domain cross-connect that is pathwise wide-sense non-blocking. Similarly F_2 has the topology of a $(2k-1) \times k$ space domain cross-connect that is pathwise wide-sense non-blocking. Then C is pathwise wide-sense non-blocking* and *wavelength strictly non-blocking.*

6 Less restrictive designs

In the previous section we presented a simple construction for a $k \times k$ strictly non-blocking WDM cross-connect. It is natural to conjecture that a more complicated or sophisticated design might be able to provide the same non-blocking characteristics with fewer than $2k-1$ wavelength interchangers. In particular the split cross-connect design required that each path pass through exactly one wavelength interchanger. While this at first may appear to be too restrictive, we will show in this section that the existence of an arbitrary $k \times k$ strictly non-blocking WDM cross-connect with fewer than $2k-1$ wavelength interchangers implies the existence of a $k' \times k'$ WDM split cross-connect that is also strictly non-blocking for some $k' < k$. This implies that the WDM split cross-connect design is as powerful as other less restrictive designs.

Let L be the set of strictly non-blocking WDM cross-connects that contain at least one directed path P from some input node $a \in I$ through $w > 1$ wavelength interchangers to some output node $b \in O$.

LEMMA 6.1. *There does not exist a 2×2 WDM cross-connect $C \in L$ with fewer than three wavelength interchangers.*

Proof. Note that any 2×2 strictly non-blocking WDM cross-connect C must have at least two wavelength interchangers since there could be two demands both with input wavelength λ_1 and output wavelength λ_2 and so these two demands could not use the same wavelength interchanger. We now show that such a cross-connect in L must in fact have at least three wavelength interchangers. By contradiction assume that there is a 2×2 cross-connect $C \in L$ with exactly two wavelength interchangers. By definition, C is strictly non-blocking and there is a directed path P in C from input node a to output node b that passes through both wavelength interchangers in C. Let e and f be the other input and output nodes respectively and let n be the number of wavelengths available. Consider what happens if we route demands $(a, \lambda_i, b, \lambda_i)$ for $1 \leq i \leq n$ along P with constant wavelength assignment λ_i. Then since P passes through all available wavelength interchangers there are no unused wavelengths in to or out of any wavelength interchanger. Thus any new demand, say $(e, \lambda_1, f, \lambda_2)$, can not possibly be satisfied by C. This contradicts the assumption that C is strictly non-blocking. ∎

LEMMA 6.2. *If for some $k > 2$, there exists a $k \times k$ WDM cross-connect $C \in L$ that has fewer than $2k-1$ wavelength interchangers, then for some $k' < k$ there exists a strictly non-blocking $k' \times k'$ WDM cross-connect $C' \notin L$ that has fewer than $2k'-1$ wavelength interchangers.*

Proof. Suppose there is some $k > 2$ for which there exists a $C \in L$ of size $k \times k$ with $m < 2k-1$ wavelength interchangers. Let n be the number of wavelengths. Since $C \in L$ we know that it is strictly non-blocking and there exists a directed path P in C from some input node $a \in I$ to some output node $b \in O$ that passes through $w > 1$ wavelength interchangers. Suppose we perform the operation $\mathbf{Fill}(C, P, n)$ that routes demands $(a, \lambda_i, b, \lambda_i)$ for $1 \leq i \leq n$ along P with constant wavelength assignment λ_i. After performing $\mathbf{Fill}(C, P, n)$, no other demand can be routed through any wavelength interchanger on P since all wavelengths are used on fibers going in to or out of such wavelength interchangers. Furthermore, all other demands that are routed through any switch s along P must enter and exit s on ingoing and outgoing fibers respectively, that are not a part of P.

Thus consider the cross-connect C' obtained by the process $\mathbf{Modify}(C, P)$ that is defined as follows. Remove input node a and output node b from C. Remove all fibers along path P. All wavelength interchangers along P are isolated (i.e. have no incoming or outgoing fibers) and so they are also removed. This construction can easily be seen to have the property that after performing **Fill**(C, P, n), any other demand will have a routing and wavelength assignment in C if and only if it does in C'.

Therefore the fact that C is strictly non-blocking implies that C' must also be strictly non-blocking. Notice that C' has size $k' \times k'$, where $k' = k - 1$, and the number of wavelength interchangers is $m - w < 2k' - 1$. Therefore C' is a strictly non-blocking cross-connect of size $k' \times k'$ with fewer than $2k' - 1$ wavelength interchangers. Notice that as long as C' contains at least one path P from some input node $a' \in I'$ through $w' > 1$ wavelength interchangers to some output node $b' \in O'$ then by definition $C' \in L$ and we can repeat this process. The size of C' will decrease by one each time **Fill**(C, P, n) and **Modify**(C, P) are performed on the current cross-connect C. By Lemma 6.1 and the fact that C' is always strictly non-blocking, eventually **Modify**(C, P) must return a strictly non-blocking WDM cross-connect $C' \notin L$ of size $k' \times k'$, where $2 \leq k' < k$, with fewer than $2k' - 1$ wavelength interchangers. ■

7 A lower bound on the number of wavelength interchangers

Our goal in this section is to show that any $k \times k$ WDM split cross-connect that is strictly non-blocking must have at least $2k - 1$ wavelength interchangers. If we consider the set of WDM cross-connects not in L and we can show that for this restricted set no strictly non-blocking $k \times k$ WDM cross-connect exists that uses fewer than $2k - 1$ wavelength interchangers, then Lemma 6.2 implies that no arbitrary strictly non-blocking WDM cross-connect of size $k \times k$ uses less than $2k - 1$ wavelength interchangers. Notice that the set of WDM cross-connects not in L includes all split cross-connects but also includes cross-connect designs that contain paths directly from some input fiber to some output fiber without passing through any wavelength interchangers.

Consider some strictly non-blocking WDM cross-connect $C \notin L$ that has $m \leq 2k - 2$ wavelength interchangers. Let λ_1 and λ_2 be two of the available wavelengths in C. For any set of demands D on C with routing R, let A_{DR} be the set of wavelength interchangers that service either a demand with input wavelength λ_1 and/or output wavelength λ_2. Let B_{DR} be the set of all other wavelength interchangers. Therefore for any demand set D and any routing R of D, $|A_{DR}| + |B_{DR}| = m$.

Given C we show that a set of valid demands and a corresponding valid routing exist that require more than $2k - 2$ wavelength interchangers if C is strictly non-blocking . For a $k \times k$ cross-connect, define a *full-*$\{\lambda_1, \lambda_2\}$ set of demands to be a valid set of $2k$ demands each of whose input and output wavelengths are in the set $\{\lambda_1, \lambda_2\}$. Notice that there exist full-$\{\lambda_1, \lambda_2\}$ sets for which any valid routing is such that each route in the routing passes through a wavelength interchanger. This follows since a full-$\{\lambda_1, \lambda_2\}$ set D can be chosen so that each $d \in D$ uses input wavelength λ_1 if it uses output wavelength λ_2 and vice versa and so any valid routing R for D must route each demand through a wavelength interchanger. In what follows, let D be such a full-$\{\lambda_1, \lambda_2\}$ set of demands and let R be a valid routing of D where necessarily each route of R passes through one wavelength interchanger. Since all demands use one of two input wavelengths and one of two output wavelengths, no wavelength interchanger can service more than two demands. Since at most $m \leq 2k - 2$ wavelength interchangers are used to service the $2k$ demands in D, at least two wavelength interchangers service two demands each. Figure 5 shows one such wavelength interchanger. Let $a_1 \in I$ and $b_1 \in O$ be the input and output fibers for one of the two demands and let $a_2 \in I$ and $b_2 \in O$ be the input and output fibers corresponding to the other demand.

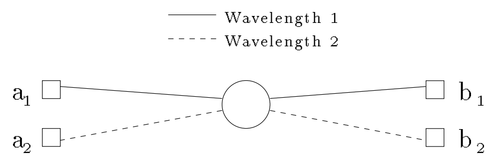

Figure 5: A wavelength interchanger that services two demands.

Suppose D and R are such that there are two demands $d_1 = (a_1, \lambda_1, b_2, \lambda_2)$ and $d_2 = (a_2, \lambda_2, b_1, \lambda_1)$ that are routed through the same wavelength interchanger WI_i. We define the operation **Uncross**(WI_i) that changes these two demands to be $(a_1, \lambda_1, b_1, \lambda_1)$ and $(a_2, \lambda_2, b_2, \lambda_2)$ and routes these new demands exactly as d_1 and d_2 were routed while keeping all other demands and routes unchanged. In other words, after **Uncross**(WI_i) is performed all fibers will have exactly the same wavelengths in use as before, every input and output node will still have two demands and each wavelength interchanger will have the same set of incoming and outgoing wavelengths.

Now we show that we can iteratively change the set D of demands and routing R of the resulting set of demands so that eventually we will have some new full-$\{\lambda_1, \lambda_2\}$ set of demands D' with a valid routing R' where all wavelength interchangers are in $A_{D'R'}$. We begin by defining a procedure **Construct**$(C,(D,R))$ that takes a strictly non-blocking $k \times k$ WDM cross-connect C not in L, a full-$\{\lambda_1,\lambda_2\}$ set of demands D and a valid routing R of D where each route in R passes through a wavelength interchanger and produces a full-$\{\lambda_1,\lambda_2\}$ set of demands D' and a valid routing R' of D' such that $|A_{D'R'}| = |A_{DR}| + 1$.

Construct(C, (D, R))

1. Take two wavelength interchangers, WI_j and WI_i, that each service two demands.
2. **Uncross**(WI_j) and **Uncross**(WI_i)
3. Let $(a_{j1},\lambda_1,b_{j1},\lambda_1)$ and $(a_{j2},\lambda_2,b_{j2},\lambda_2)$ be the two resulting demands that WI_j services.
4. Let $(a_{i1},\lambda_1,b_{i1},\lambda_1)$ and $(a_{i2},\lambda_2,b_{i2},\lambda_2)$ be the two resulting demands that WI_i services.
5. Remove $(a_{j1},\lambda_1,b_{j1},\lambda_1)$ and $(a_{i2},\lambda_2,b_{i2},\lambda_2)$ from D and route all remaining demands according to R to create D^* and R^*.
6. Add $d_1 = (a_{j1},\lambda_1,b_{i2},\lambda_2)$ and $d_2 = (a_{i2},\lambda_2,b_{j1},\lambda_1)$ to D^* to create D'.
7. Route all demands in D' that are also in D^* according to R^*. Add a valid route for each of the two new demands d_1 and d_2 to R^* to create R'.
8. Return (D',R')

First note that by assumption C is strictly non-blocking and therefore Step 7 of **Construct**$(C,(D,R))$ must always be possible. Figure 6 shows WI_j and WI_i first under the original set D of demands and then under the new set D' of demands. Note that WI_x and WI_z may be the same wavelength interchanger. Also, in order for Step 1 to always be possible it must be that each route in R passes through one wavelength interchanger. As mentioned earlier, this can be achieved initially by choosing the original D so that each demand in D has input wavelength that differs from its output wavelength. Then it should be noticed that Step 7 must route d_1 and d_2 through a wavelength interchanger and so the resulting routing R' will again have the property that all routes in R' pass through one wavelength translator.

LEMMA 7.1. *If D is a full-$\{\lambda_1,\lambda_2\}$ set of demands and R is a valid routing in C that satisfies D, then (D',R') =* **Construct**$(C,(D,R))$ *is such that D' is a full-$\{\lambda_1,\lambda_2\}$ set of demands, R' is a valid routing for D' and $|A_{D'R'}| = |A_{DR}| + 1$.*

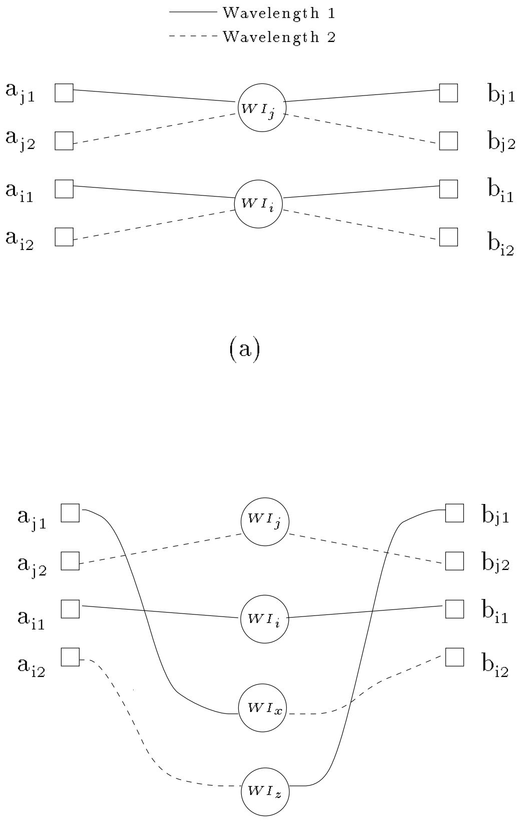

(a)

(b)

Figure 6: In **a)** WI_j and WI_i are each servicing two demands in order for C to satisfy D. In **b)** WI_j and WI_i and two other wavelength interchangers under the new set D' of demands on C.

Proof. Clearly D' is a full-$\{\lambda_1,\lambda_2\}$ set of demands. Consider WI_a, a wavelength interchanger that is in A_{DR}. Suppose that the set of demands that WI_a services in R' includes the set of demands that it serviced in R. Then WI_a is also in $A_{D'R'}$. The only wavelength interchangers in A_{DR} that service demands in R that they do not service in R' are the two wavelength interchangers, WI_j and WI_i, that service two demands in R. Since in R' they still each service a demand with either input wavelength λ_1 or output wavelength λ_2, it must be that they are both in $A_{D'R'}$.

Consider first the wavelength interchanger WI_z that must service the new demand d_2 of Step 6. Note that routing d_2 through WI_z will not cause WI_z to service a demand that has an input wavelength λ_1 or an output wavelength λ_2 if WI_z serviced no such demand in R. Therefore the route chosen for d_2 does not increase or decrease the size of $A_{D'R'}$ with respect to the size of A_{DR}.

Now consider WI_x, the wavelength interchanger that services the other new demand d_1 in Step 6. By definition all wavelength interchangers in A_{DR} service a demand with input wavelength λ_1 and/or output wavelength λ_2. This implies that R' can route d_1 through WI_x if and only if $WI_x \notin A_{DR}$. Since d_1 has input wavelength λ_1 (and output wavelength λ_2), routing d_1 through WI_x in R' implies $WI_x \in A_{D'R'}$. Thus $|A_{D'R'}| = |A_{DR}| + 1$ after C has satisfied all demands in D' using R'. ■

Notice that we can use **Construct**$(C, (D, R))$ as long as there are at least two wavelength interchangers in C that service two demands. By definition **Construct**$(C, (D, R))$ creates a new full-$\{\lambda_1, \lambda_2\}$ set of demands D'. Given that there is no route in R that fails to pass through any wavelength interchanger, we know that the same is true for the routing R'. Since m, the number of wavelength interchangers in C, is less than $2k-1$, C must always use at least two wavelength interchangers to service two demands each for any set D' of demands and R' of routes that **Construct**$(C, (D, R))$ returns. Therefore we can always call **Construct**$(C, (D, R))$ on the current set of demands and routes. Furthermore since $|B_{DR}| = m - |A_{DR}|$ and since **Construct**$(C, (D, R))$ returns a set D' of demands and R' routes for which $|A_{D'R'}| = |A_{DR}| + 1$, **Construct**$(C, (D, R))$ also returns a set D' of demands and R' routes for which $|B_{D'R'}| = |B_{DR}| - 1$. Therefore we can call **Construct**$(C, (D, R))$ repeatedly until $|B_{DR}| = 0$. We have now shown that it is possible to create a full-$\{\lambda_1, \lambda_2\}$ set of demands, D, that can be routed in such a way that all wavelength interchangers in C are also in A_{DR}. Consider what happens if we then make one more call to **Construct**$(C, (D, R))$. Since C is strictly non-blocking it must be able to service both of the new demands d_1 and d_2 in Step 6. However, since no wavelength interchanger in A_{DR} can service d_1 and since $|B_{DR}| = 0$, C is not strictly non-blocking. Since this holds for any $m \le 2k-2$, any strictly non-blocking WDM split cross-connect must have at least $2k-1$ wavelength interchangers.

LEMMA 7.2. *Suppose $C \notin L$ is a WDM cross-connect with k input fibers, k output fibers and $n > 1$ available wavelengths. Then C must have at least $2k-1$ wavelength interchangers if it is pathwise and wavelength strictly non-blocking.*

Proof. The result follows from the discussion above. ■

THEOREM 7.3. *If any WDM cross-connect with k input fibers, k output fibers and $n > 1$ available wavelengths is pathwise and wavelength strictly non-blocking it must have at least $2k-1$ wavelength interchangers.*

Proof. Follows directly from Lemma 6.2 and Lemma 7.2. ■

8 Discussion

In [KK99, WW98] another problem concerned with minimizing the number of wavelength interchangers is considered. In these papers, an optical network is given and the goal is to determine the minimum number of nodes in the network such that if wavelength interchange is allowed at these nodes the resulting network needs no more wavelengths to route any set of demands than if wavelength interchange was allowed at every node. For this problem, it was assumed that a single wavelength interchanger at a node would provide complete wavelength interchange capability regardless of the number of fibers entering and exiting the node. Also, there was no concern with having the resulting network satisfy strictly non-blocking conditions. However, if the non-blocking constraint is added then the lower bound result presented in Section 7 implies that if a node with indegree and outdegree of k is chosen as one that will have wavelength interchange capability then at least $2k-1$ wavelength interchangers will be needed at that node. Thus the number of wavelength interchangers used will be proportional to the sum of the degrees of the nodes chosen rather than the number of nodes.

Acknowledgments

We are very grateful to Bruce Shepherd for his significant ongoing contributions that assisted greatly in the development of this work. We also thank Tim Griffin for his many helpful comments.

References

[Ber73] C. Berge. *Graphs and Hypergraphs.* North Holland, Amsterdam, 1973.

[CH82] R. Cole and J. E. Hopcroft. On edge coloring bipartite graphs. *SIAM J. Comput.*, 11(3):540–546, 1982.

[DMR+99] C. R. Doerr, B. Mikkelsen, G. Raybon, P. Schiffer, L. W. Stultz, M. Zirngibl, G. Wilfong, M. Cappuzzo, E. Laskowski, A. Paunescu, L. Gomez, and J. Gates. Wavelength selective cross connect using arrayed waveguide lens multi-wavelength filters. In *Proc. of OFC'99*, 1999.

[Gab76] H. N. Gabow. Using Euler partitions to edge color bipartite graphs. *International Journal of Computer and Information Sciences*, 5(4):345–355, 1976.

[GK82] H. N. Gabow and O. Kariv. Algorithms for edge coloring bipartite graphs and multigraphs. *SIAM J. Comput.*, 11(1):117–129, 1982.

[Hui90] J. Y. Hui. *Switching and Traffic Theory for Integrated Broadband Networks*. Kluwer Academic Publishers, Norwell, MA, 1990.

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

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