Strictly Non-blocking WDM Cross-connects

by

April Rasala
B. A.
Dartmouth College, 1999

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Master of Science in Electric Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 2001

Author ....
Department of Electrical Engineering and Computer Science
May 7, 2001

Certified by ....
David R. Karger
Associate Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by ....
Arthur C. Smith
Chairman, Departmental Committee on Graduate Students

Strictly Non-blocking WDM Cross-connects

by

April Rasala

Submitted to the Department of Electrical Engineering and Computer Science on May 7, 2001, in partial fulfillment of the requirements for the degree of Master of Science in Electric Engineering and Computer Science

Abstract

Using wavelength division multiplexing (WDM) technology, an optical network can route multiple signals simultaneously along a single optical fiber by encoding each signal on its own wavelength. If the network contains places where multiple fibers connect together and signals are allowed to be moved from any of the incoming fibers to any of the outgoing fibers, then the network is said to contain *cross-connects*. More precisely, a $k_1 \times k_2$ WDM *cross-connect* has k_1 input fibers and k_2 output fibers. Each of the k_1 input fibers supports the same n_1 input wavelengths and each of the k_2 output fibers supports the same n_2 output wavelengths. Since a signal on input wavelength λ can be routed from its input fiber to an output fiber such that it arrives on the output fiber using wavelength γ, where $\lambda \neq \gamma$, the cross-connect must be capable of performing wavelength conversion. Along any fiber in the cross-connect a device called a *wavelength interchanger* can be inserted to perform wavelength conversion. In other words if the path of a signal from an input fiber to an output fiber passes through a wavelength interchanger, then the wavelength of the signal can be changed to any wavelength that is not already in use along the fiber leaving the wavelength interchanger. Given the high cost of wavelength interchangers, the overall cost of a $k_1 \times k_2$ WDM cross-connect is minimized by reducing the number of wavelength interchangers in the cross-connect. However, a desirable property for a cross-connect C is for C to always be able to provide a route (and wavelength conversion) for any valid demand from any pair of input and output fibers regardless of the routes of other demands currently routed in C. If C has this capability then it is said to be *strictly non-blocking*.

For most of this thesis we consider a demand to be a request for a connection from an input fiber to an output fiber such that the connection starts on a specified input wavelength and leaves the cross-connect on a second specified wavelength. Using this demand model, we consider cross-connects for which k_1 is not necessarily equal to k_2 and the number n_1 of supported input wavelengths can differ from the number n_2 of supported output wavelengths. Without loss of generality we assume that $k_1 \leq k_2$ and present a family of strictly non-blocking $k_1 \times k_2$ WDM cross-connects that use $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. For the case when $k_1 = k_2 = k$ and $n_1 = n_2$ we prove that this is optimal. For cross-connects where n_1 is not necessarily equal to n_2, we show that if there is at most one wavelength interchanger on any path from an input fiber to an output fiber, $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers are optimal. Finally, we consider a more flexible demand model where $k_1 = k_2$ but the input and output wavelengths are not specified as part of the demand. We show that $2k - 1$ wavelength interchangers are still necessary for any strictly non-blocking $k \times k$ WDM cross-connect.

Thesis Supervisor: David R. Karger

Title: Associate Professor of Electrical Engineering and Computer Science

Acknowledgments

The problem considered in this thesis was first posed to me by Gordon Wilfong. I wish to thank him for his contribution to the results presented here and his advice throughout this process. I wish to thank David Karger for his support and constructive comments. I would like to thank Richard Rasala for his helpful feedback. Finally, I would like to thank the many students who have encouraged me along the way.

Contents

Chapter 1

Introduction

A $k_1 \times k_2$ wavelength division multiplexing (WDM) cross-connect is a directed network of fibers and optical devices that allows for signals to be moved from the set of k_1 input fibers to the set of k_2 output fibers. Each input fiber can carry up to n_1 signals simultaneously by encoding each signal on its own wavelength. Similarly each output fiber can carry up to n_2 signals simultaneously. We restrict our attention to wavelength interchanging WDM cross-connects. In such cross-connects the wavelength on which the signal enters the cross-connect can be changed at a wavelength interchanger. A wavelength interchanger is an optical device with a single input fiber and a single output fiber. Any of the signals entering the wavelength interchanger on any of the supported wavelengths can be changed to any of the supported wavelengths on the fiber leaving the wavelength interchanger provided those wavelengths are not already in use by another connection. The flexibility to change the wavelength of a signal passing through the cross-connect allows demands on the cross-connect to specify the input fiber, input wavelength, output fiber and output wavelength of any connection request. This provides for the greatest level of flexibility for the network using the cross-connect to move signals from a set of fibers to another set of fibers (and possibly another set of wavelengths). Throughout this thesis we will refer to a WDM cross-connect and mean a wavelength interchanging WDM cross-connect.

A WDM cross-connect must maintain current connections (the results of previous demands placed on it) and at the same time satisfy new requests for connections between input fibers and output fibers. A *strictly non-blocking* cross-connect C is guaranteed to be able to route any new demand placed on it regardless of the set of existing connections routed through C.

Given the high cost of wavelength interchangers now and in the foreseeable future [14], the cost of a strictly non-blocking WDM cross-connects is dominated by the number of wavelength interchangers required in the design. Thus we focus our attention on determining the optimal number of wavelength interchangers needed in any strictly non-blocking $k_1 \times k_2$ WDM cross-connect and presenting a family of strictly non-blocking $k_1 \times k_2$ WDM cross-connects that uses the optimal number of wavelength interchangers.

The task of optimizing strictly non-blocking WDM cross-connects is simplified by drawing from work on classical cross-connects. In classical cross-connects, signals do not share wires and therefore there is no need for wavelength conversion. A vast literature exists on the design of such cross-connects [4, 8, 13, 1, 7]. In particular, work has been done on the optimization of strictly non-blocking classical cross-connects that guarantee that a connection can always be made between any input line and any output line provided neither are already in use. The topology of a strictly non-blocking classical cross-connect can be used as the topology of a WDM network. In this case we assume that the WDM network supports n wavelength and does not allow any wavelength conversion. Since the WDM network has the topology of a classical network and it does not support wavelength conversion, it is conceptually equivalent to n copies, one for each wavelength, of the classical network. Furthermore, since the WDM network has the topology of a strictly non-blocking classical cross-connect, any connection on any wavelength can be completed as long as the requesting input and output fibers do not already have an existing connection on that wavelength. We refer to a section of directed fiber and optical devices without wavelength conversion capability as a *fabric*.

The designs of strictly non-blocking WDM cross-connects presented in this paper use the topology of classical strictly non-blocking cross-connects as the basis for fabrics that connect the input fibers to a set of wavelength interchangers and connect those wavelength interchangers to the set of output fibers. Such a design is referred to as a split cross-connect. For the case where the number of input fibers, k, is equal to the number of output fibers and the set of input wavelengths is the same as the set of output wavelengths, we present split cross-connects with $2k-1$ wavelength interchangers and prove that these designs are optimal. We also consider general split cross-connects where the number of input fibers, k_1, and the number of input wavelengths, n_1, may differ from the number of output fibers, k_2, and the number of output wavelengths, n_2. In this case we present designs of strictly non-blocking split cross-connects that use $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. We show that this is optimal for the class of split cross-connects. However, we leave as an open question whether the family of split cross-connects are optimal if the number of input fibers is not equal to the number of output fibers and the set of supported input wavelengths differs from the set of supported output wavelengths.

1.1 Definitions

1.1.1 Wavelength Division Multiplexing

Wavelength division multiplexing (WDM) technology allows multiple signals to be routed along an optical fiber simultaneously by encoding each signal on its own wavelength. For example, a WDM network might allow each fiber to carry up to n signals simultaneously by giving each signal its own wavelength out of a set of n supported wavelengths. Devices crucial to such networks are *optical switches*. An optical switch has an arbitrary number of input fibers and output fibers and any signal encoded on any wavelength on any of the input fibers can be routed to the same wavelength on any of the outgoing fibers. If these are the only devices included in an optical network, then a connection that is originally routed on a particular wavelength, say λ, must stay on that wavelength for its entire route. Even in simple networks this can result in congestion of the network even when no single fiber is carrying all n signals [6, 15]. Therefore it is often desirable to allow the wavelength of a connection to be changed. A *wavelength interchanger* is a device that has a single input fiber and a single output fiber. A wavelength interchanger can change the wavelength of any subset of the signals entering it provided the wavelengths that it encodes those signals on are unused on its output fiber. See Figure 1-1.

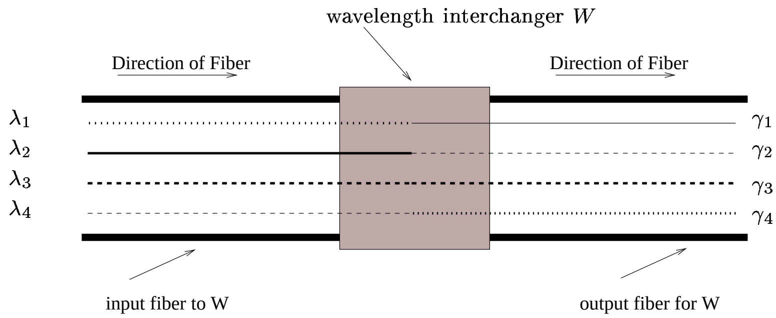

Figure 1-1: A wavelength interchanger W with four signals passing through it. The signals on input wavelengths λ_1, λ_2, λ_3 and λ_4 are changed to output wavelengths γ_1, γ_2, γ_3 and γ_4 respectively

1.1.2 WDM Cross-connects

In a WDM network, fibers join at various places and signals on the incoming fibers can be routed to the outgoing fibers. The part of the network that allows this routing of signals from some set of incoming fibers to a set of outgoing fibers is referred to as a WDM *cross-connect*. A $k_1 \times k_2$ WDM cross-connect has k_1 input fibers and k_2 output fibers. We define $I = \{I_1 \ldots I_{k_1}\}$ to be the set of input fibers and $O = \{O_1, \ldots O_{k_2}\}$ to be the set of output fibers. The set of n_1 wavelengths, $\{\lambda_1, \ldots \lambda_{n_1}\}$, supported on each of the k_1 input fibers is referred to as the set of *input wavelengths* and similarly the set of *output wavelengths* is the set of n_2 wavelengths, $\{\gamma_1, \ldots \gamma_{n_2}\}$ supported on each of the output fibers. Without loss of generality we assume that $k_1 \leq k_2$ and $n_1, n_2 \geq 2$.

In general a demand on a WDM cross-connect could be specified in various ways. For this thesis we define a *demand*, $d = (I_x, \lambda_y, O_z, \gamma_w)$ to be a request for a path and an assignment of wavelengths along the path from input fiber I_x to output fiber O_z such that the connection starts on input wavelength λ_y, changes wavelength only at a wavelength interchanger and ends on output wavelength γ_w. Such a path and wavelength assignment is called a *route*. A routing R of D is a collection of routes, one for each demand $d \in D$, such that no two demands are assigned the same wavelength along a shared section of fiber. We will show that for some classes of cross-connects, this demand and route model is equivalent to models in which a demand does not specify the input and/or the output wavelength and a route is allowed to use any wavelengths available on the specified input and output fibers. We leave as an open question whether more flexible demand models allow for improved results for other types of cross-connects.

Given an existing set of demands, D, we say demand, $d = (I_x, \lambda_y, O_z, \gamma_w)$, is *valid* with respect to D if input wavelength λ_y is unused on input fiber I_x by all demands in D and output wavelength γ_w is unused on output fiber O_z by all demands in D. A route r for d is *valid* with respect to a routing R of D if the wavelength assigned to d along each fiber in r is not used by any other connection currently routed on the fiber according to R.

1.1.3 Non-blocking Properties

Ideally a cross-connect will always be able to satisfy a valid demand with a valid route. Since this can be done in various ways, different levels of *non-blocking* have been defined as follows.

- **Rearrangably Non-blocking:** A cross-connect C is rearrangably non-blocking if any new valid demand d can be assigned a new valid route r provided that any necessary existing routes are rerouted.

- **Wide-sense Non-blocking:** A cross-connect C is wide-sense non-blocking if any new valid demand d can be assigned a route by an algorithm A provided that all currently routed demands were routed using A.

- **Strictly Non-blocking:** A cross-connect C is strictly non-blocking if any new valid demand d can be assigned a route regardless of the routes assigned to any currently routed demands.

Strictly non-blocking and wide-sense non-blocking cross-connects have many advantages over rearrangably non-blocking cross-connects. For example, a rearrangably non-blocking cross-connect will create a buffering problem when it must interrupt currently routed connections to reroute them. While both strictly non-blocking and wide-sense non-blocking cross-connects avoid this buffering problem, a strictly non-blocking cross-connect is perhaps more robust since it does not depend upon having taken care to route previously requested demands in order to guarantee a route for any new demands. This is particularly desirable in practice because often times hardware failures necessitate that some routes are chosen without regard to a particular algorithm. A wide-sense non-blocking cross-connect might require that all routes are rerouted after such a hardware failure. In contrast, since a strictly non-blocking cross-connect does not depend on the cross-connect having routed existed demands carefully, it is capable of handling situations where the set of existing routes may have been chosen to work around a hardware failure without regard for how they utilize the resources inside the cross-connect. We focus our attention on the design of strictly non-blocking $k_1 \times k_2$ WDM cross-connects but note that the design of optimal wide-sense non-blocking cross-connects would also be of interest.

1.1.4 Topologies of Cross-connects

Split Cross-connects

A *fabric* is defined to be a subnetwork that does not contain any wavelength interchangers. A $k_1 \times k_2$ WDM cross-connect is said to be a split cross-connect if it has k_1 input fibers connected to a fabric, F_1, and k_2 output fibers connected to a second fabric, F_2. Furthermore the output fibers of F_1 are the input fibers for a set W of wavelength interchangers and the output fibers for the set W of wavelength interchangers are the input fibers of the fabric F_2. Another way to say this is to say that C is a split cross-connect if every path from an input fiber to an output fiber through C has exactly one wavelength interchanger. See Figure 1-2.

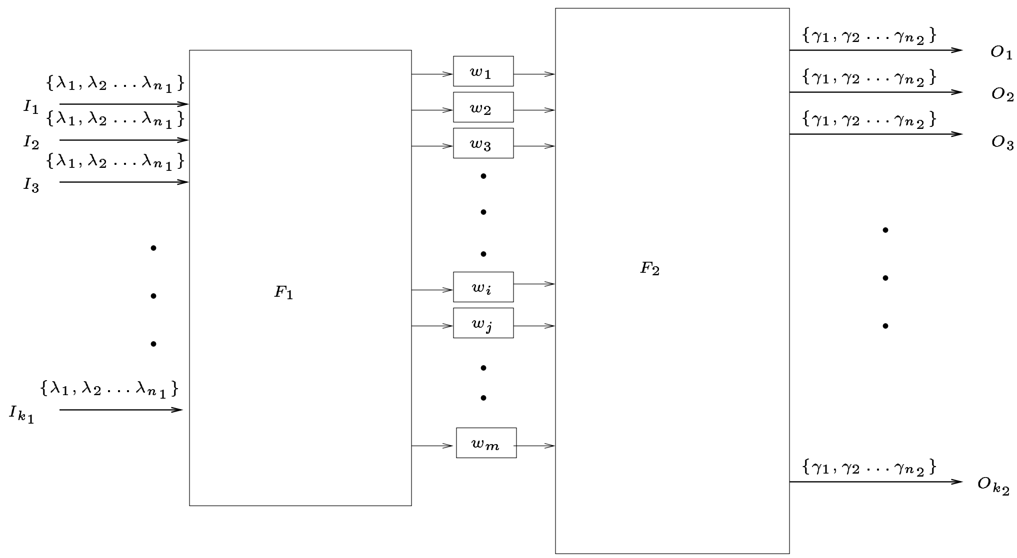

Figure 1-2: A $k_1 \times k_2$ WDM split cross-connect with m wavelength interchangers. The set of k_1 input fibers are connected to the m wavelength interchangers with a fabric F_1. The set of m wavelength interchangers are connected to the k_2 output fibers with a fabric F_2.

The designs for strictly non-blocking cross-connects that we present are all split cross-connects. We will prove that in the case where the number of input and output fibers is the same, this is optimal. Note that without loss of generality we will overload the definition of split and allow it to also refer to any cross-connect for which every path p from any input fiber to any output fiber has *at most* one wavelength interchanger.

Homogeneous and Heterogeneous Cross-connects

In this thesis we consider cross-connects for two different situations. The first situation involves cross-connects that have the same number of input fibers as output fibers. In these cross-connects the same set of input and output wavelengths are supported throughout the cross-connect. We call such cross-connects *homogeneous* cross-connects.

In the second case, the purpose of the cross-connect is to connect a set of fibers from a given network to a set of fibers from another network. In this case the set of wavelengths supported on the input fibers may differ from the set of wavelengths supported on the output fibers. Also the number of input fibers may differ from the number of output fibers. We call such a cross-connect a *heterogeneous* cross-connect.

1.2 Previous Results

This work draws both from the literature of classical cross-connects and from more recent work on WDM cross-connects.

A classical architecture considers the case where each input and output line carries only one signal at a time. Designs for cross-connects in classical networks have been well studied

[3, 13, 8, 1, 2, 7]. In a classical network the optimization criterion is the number of edges in the graph representing the topology of the cross-connect. In contrast, in wavelength interchanging WDM cross-connects the optimality criterion is the number of wavelength interchangers.

In general the cost of a design of a WDM network is dominated by the number of wavelength interchangers needed in the network. Thus the problem of optimizing WDM networks has been studied in a model that assigns a cost of one to any node in the network where wavelength conversion capability is provided [15, 6].

Other work has focused on the possibility of using less powerful (and therefore less expensive) wavelength conversion devices. One example would be a device that was capable of swapping two wavelengths while leaving the others fixed. That work has considered what effect using such less powerful wavelength interchangers has on the number of wavelength interchangers needed in total[9].

Recently some work has been done on considering the cost of providing complete wavelength conversion capability in a WDM cross-connect. For instance, various wavelength interchanging WDM homogeneous cross-connects with a variety of non-blocking capabilities have been considered and in particular one design that used $k \log k$ wavelength interchangers was shown to be strictly non-blocking [14]. We will improve upon this result by presenting a design for a strictly non-blocking cross-connect with $2k - 1$ wavelength interchangers.

1.3 Our Contribution

We will consider homogeneous and heterogeneous cross-connects. We present the design of a strictly non-blocking $k_1 \times k_2$ WDM cross-connect that uses $k_1 + k_2 - 1$ wavelength interchangers. We also present the design of strictly non-blocking $k_1 \times k_2$ WDM cross-connects that use $n_1 k_1$ wavelength interchangers. Together these provide a family of strictly non-blocking $k_1 \times k_2$ WDM cross-connects that use at most $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers.

When considering the optimality of these designs we handle the homogeneous and heterogeneous cases separately. We start our discussion of the optimal number of wavelength interchangers in a homogeneous cross-connect by showing that if there exists a family of optimal non-split cross-connects, then there exists a family of split cross-connects that use the same number of wavelength interchangers. This reduces the problem to a special case of showing that a heterogeneous $k \times k$ WDM split cross-connect requires $2k-1$ wavelength interchangers. For the heterogeneous case we consider only split cross-connects and leave as an open question whether there exist non-split cross-connects with strictly fewer wavelength interchangers. We present a sequence of demands that requires that any strictly non-blocking $k_1 \times k_2$ WDM heterogeneous split cross-connect use at least $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. These results together show that the family of strictly non-blocking $k_1 \times k_2$ WDM cross-connects presented in Chapter 2 are optimal for homogeneous cross-connects and heterogeneous split cross-connects. Finally, we consider more relaxed demand models where the input and output wavelengths are not specified as part of the demand. We show that even under this model of demands, $2k-1$ wavelength interchangers are necessary and sufficient for homogeneous cross-connects. The work presented in this thesis is joint work with Gordon Wilfong [12, 10].

1.3.1 Overview

In Chapter 2 we present the design of a family of $k_1 \times k_2$ WDM cross-connects. We show that these are strictly non-blocking. For the case of homogeneous cross-connects we prove, in Chapter 3, that if there is a family of optimal non-split cross-connects, then there is a family of optimal split cross-connects with the same number of wavelength interchangers. We then restrict our attention to cross-connects with split designs. We present a sequence of demands that shows that any homogeneous strictly non-blocking $k \times k$ WDM cross-connect must have $2k - 1$ wavelength interchangers. We then show that the lower bounds presented hold under demand models that do not specify the input and output wavelength as part of the demand. For the heterogeneous case we show, in Chapter 4, that any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect must have $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. Finally in Chapter 5 we consider future work.

Chapter 2

Designs for a Family of Strictly Non-blocking Cross-connects

2.1 Traditional Cross-connects

We start our discussion with a look at traditional cross-connect designs and how they apply to the design of WDM cross-connects. In a classical network, a *line* carried only one *channel* (wavelength) and therefore no two signals shared a line. A traditional strictly non-blocking $k_1 \times k_2$ cross-connect had k_1 input lines and k_2 output lines. A demand was a request for a connection from an input line to an output line. Note that in this setting at most one demand could use any particular input or output line at a given time. A route for a demand in a classical cross-connect was a path from the input line to the output line that was edge-disjoint from any existing route in the cross-connect. Therefore we say that the topology of a strictly non-blocking classical cross-connect is pathwise strictly non-blocking.

In the WDM setting we require that no two signals on the same wavelength share a fiber. Therefore if a WDM fabric has no wavelength conversion capability, then a demand will be a request for a connection on a specified wavelength from an input fiber to an output fiber. A route for such a demand is a path from the input fiber to the output fiber that is edge-disjoint (fiber-disjoint) from all other existing demands on the same wavelength. Therefore it is natural to extend the designs of classical cross-connects to WDM fabrics as follows.

Suppose we consider the topology of a classical cross-connect C. The topology of C can be represented by a graph, G, where the edges correspond to lines in C and the nodes correspond to components in C. We create a WDM fabric F from G by putting a WDM fiber in F, supporting n wavelengths, for each edge in G and a WDM switch in F for each node in G. Notice that since we have not placed any wavelength interchangers in F a signal passing through F must stay on the same wavelength for its entire path through F. Furthermore since signals on different wavelengths do not interfere with each other in WDM fabrics, one can conceptually think of F as allowing us to have n copies of C, one for each of the n wavelengths. Therefore if C was pathwise strictly non-blocking, then for each wavelength λ_i we are guaranteed that we can route a demand from input fiber a to output fiber b on wavelength λ_i as long as both a and b currently do not have a demand that uses wavelength λ_i.

A WDM fabric that is pathwise strictly non-blocking greatly simplifies the design of a strictly non-blocking WDM cross-connect and therefore we will make use of this idea when presenting our designs.

2.2 Split Cross-connects

We start with the design of a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect. First we connect the input fibers to a set of $k_1 + k_2 - 1$ wavelength interchangers with a fabric F_1. The only restriction on F_1 is that it must have the topology of a pathwise strictly non-blocking $k_1 \times (k_1 + k_2 - 1)$ classical cross-connect. We then connect the set of $k_1 + k_2 - 1$ wavelength interchangers to the input side of another fabric F_2. The output side of F_2 is the set of k_2 output fibers of our cross-connect. The only restriction on F_2 is that it has the topology of a pathwise strictly non-blocking $(k_1 + k_2 - 1) \times k_2$ classical cross-connect. This family of strictly non-blocking $k_1 \times k_2$ WDM cross-connects is by definition a family of split cross-connects. See Figure 2-1.

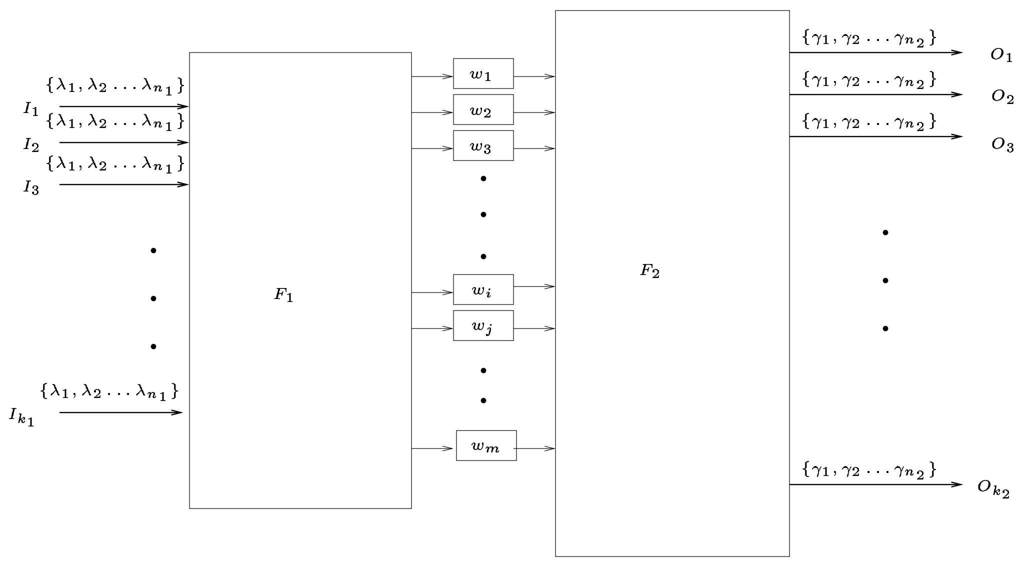

Figure 2-1: A $k_1 \times k_2$ WDM split cross-connect with $m = k_1 + k_2 - 1$ wavelength interchangers. The set of k_1 input fibers are connected to the m wavelength interchangers with a fabric F_1. The set of m wavelength interchangers are connected to the k_2 output fibers with a fabric F_2.

Theorem 1 *Let C be a $k_1 \times k_2$ WDM split cross-connect where the number of wavelength interchangers is $k_1 + k_2 - 1$. Suppose that the topology of F_1 is that of some pathwise strictly non-blocking $k_1 \times (k_1 + k_2 - 1)$ cross-connect. Similarly F_2 has the topology of a pathwise strictly non-blocking $(k_1 + k_2 - 1) \times k_2$ cross-connect. Then C is a strictly non-blocking $k_1 \times k_2$ WDM cross-connect.*

Proof: To show that C is strictly non-blocking it is enough to show that for any existing set of demands D that are currently routed through C and any new valid demand $d = (a, \lambda_1, b, \gamma_2)$ where $D \bigcup d$ is a valid demand set, there is a valid route for d. Let R be any valid routing of D

We will break D into three sets of demands. First let D_1 be the set of demands in D that use λ_1 as their input wavelength and let D_2 be the set of demands in D that use γ_2 as their output wavelength. Finally let D_3 be $D \setminus (D_1 \bigcup D_2)$. Notice that no demands in D_3 can possibly keep us from routing d since these demands all use a wavelength other than λ_1 as they pass through F_1 and a wavelength other than γ_2 as they pass through F_2. Therefore we restrict our attention to showing that d cannot be blocked by demands in $D_1 \bigcup D_2$.

Let W be the set of wavelength interchangers without a demand in D_1 or D_2 routed through them. First we show that W is non-empty. Since d is a valid demand, input fiber a could not have a demand in D that used input wavelength λ_1 and output fiber b could not have a demand in D that used output wavelength γ_2. Hence, $|D_1| \leq k_1 - 1$ and $|D_2 \leq k_2 - 1$ and therefore $|D_1 \bigcup D_2| \leq k_1 + k_2 - 2$. Since C is a split cross-connect each route passes through at most one wavelength interchanger. Therefore $|W| \geq k_1 + k_2 - 1 - |D_1 \bigcup D_2| \geq 1$.

Let $w \in W$ be a wavelength interchanger that does not currently have a demand routed through it using either λ_1 or γ_2. If there exists a route from a to w and a route from w to b, then there exists a route for d through C. However, since F_1 has the topology of a pathwise strictly non-blocking $k_1 \times (k_1 + k_2 - 1)$ cross-connect, there must exist a path from a to w that is edge-disjoint from all other paths for demands in D_1. Similarly there must exist a path in F_2 from w to b that is edge-disjoint from all other paths through F_2 for demands in D_2.

Therefore there exists a valid route through C for d and as a result C is strictly non-blocking. ■

We call this type of strictly non-blocking $k_1 \times k_2$ WDM cross-connect, a *standard* strictly non-blocking $k_1 \times k_2$ WDM cross-connect.

2.3 Dedicated Cross-connects

Note that the family of standard $k_1 \times k_2$ WDM cross-connects presented in the previous section uses $k_1 + k_2 - 1$ wavelength interchangers. For homogeneous cross-connects we will show in Chapter 3 that this is optimal. However for heterogeneous cross-connects we must also consider a limiting case.

Notice that for a heterogeneous $k_1 \times k_2$ WDM cross-connect with n_1 supported input wavelengths on each input fiber, the maximum number of demands in any valid demand set is $n_1 k_1$. In particular, if $k_2 > (n_1 - 1)k_1$ then the maximum number of possible demands is no more than $k_1 + k_2 - 1$. In this case we will use what we refer to as a dedicated cross-connect.

In a *dedicated* $k_1 \times k_2$ WDM cross-connect each input fiber is connected to its own $1 \times n_1$ optical switch and the n_1 output fibers of this switch are then connected to n_1 wavelength interchangers. Conceptually each input fiber has n_1 wavelength interchangers dedicated to service demands from it. The set of $n_1 k_1$ wavelength interchangers are connected to the set of output fibers with a fabric F_2 that has the topology of a pathwise strictly non-blocking $(n_1 k_1) \times k_2$ cross-connect. See Figure 2-2.

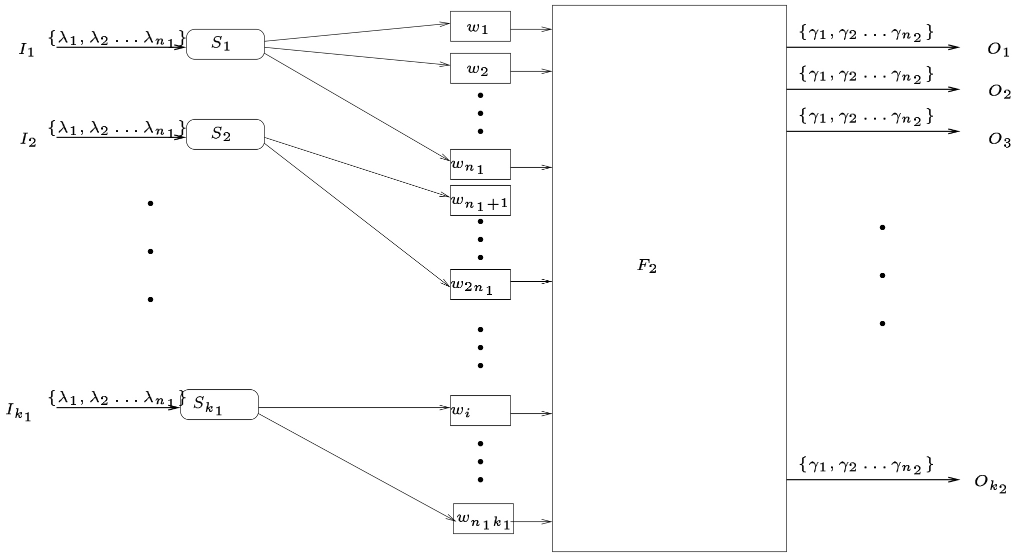

Figure 2-2: A $k_1 \times k_2$ WDM dedicated cross-connect with $n_1 k_1$ wavelength interchangers. For each of the k_1 input fibers, I_i there is a $1 \times n_1$ switch S_i. The n_1 output fibers of S_i are connected to n_1 wavelength interchangers that are used only for demands from I_i. The set of $n_1 k_1 +$ wavelength interchangers are connected to the k_2 output fibers with a fabric F_2.

Theorem 2 *Let C be a $k_1 \times k_2$ WDM dedicated cross-connect where the number of wavelength interchangers is $n_1 k_1$. Then C is strictly non-blocking.*

Proof: There is one wavelength interchanger reserved for each possible demand and F_2 has the topology of a pathwise strictly non-blocking cross-connect.

2.4 Minimal Strictly Non-blocking Cross-connects

By choosing either a dedicated or standard cross-connect this set of designs provides a family of $k_1 \times k_2$ WDM cross-connects that use $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. In the next chapter we prove that this is optimal for homogeneous cross-connects and heterogeneous split cross-connects.

Chapter 3

Lower Bound for Homogeneous Cross-connects

In Chapter 2 we presented a family of strictly non-blocking $k_1 \times k_2$ WDM cross-connects that use $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers. If the cross-connect is a homogeneous cross-connect with k input fibers and k output fibers, these designs use $2k - 1$ wavelength interchangers. In this chapter we present matching lower bounds. In general the flavor of these proofs is that we assume that we have a strictly non-blocking $k \times k$ WDM cross-connect that uses $m < 2k - 1$ wavelength interchangers and show that there is a sequence of demands and a routing of those demands such that there is a final valid demand d that cannot be routed by the cross-connect.

Although we presented our upper bounds for homogeneous and heterogeneous cases together, in this chapter we handle only the homogeneous case and assume that $k = k_1 = k_2$, $n = n_1 = n_2$ and the set of input wavelengths, $\{\lambda_1 \dots \lambda_n\}$, is also the set of output wavelengths. We leave the heterogeneous case for the next chapter.

Initially, we prove our lower bounds using a demand model where the input and output wavelengths are specified as part of the demand. In this setting we prove that if there is some $k \geq 3$ for which there is a strictly non-blocking $k \times k$ WDM cross-connect with $m < 2k - 1$ wavelength interchangers, then there exists some $k' \geq 2$ for which there is a strictly non-blocking $k' \times k'$ WDM *split* cross-connect with $m' < 2k' - 1$ wavelength interchangers. We then show that any strictly non-blocking $k \times k$ WDM *split* cross-connect requires $2k - 1$ wavelength interchangers. Combining these results shows that any strictly non-blocking $k \times k$ WDM cross-connect requires $2k - 1$ wavelength interchangers.

We then consider a demand model in which a demand consists of only an input and output fiber and the cross-connect is allowed to choose the input and output wavelengths from the set of available input and output wavelengths. We present a simple construction that extends the proofs in this chapter to hold in this new demand model.

3.1 The Adversarial View of a Strictly Non-blocking Cross-connect

In this chapter we will be proving a lower bound on the number of wavelength interchangers necessary in any strictly non-blocking $k \times k$ WDM cross-connect. To prove such a lower bound, we will present a sequence of demands and a sequence of valid routings the cross-connect could have chosen that eventually force the cross-connect to use at least $2k-1$ wavelength interchangers.

We assume that we are given a strictly non-blocking $k \times k$ WDM cross-connect C. As an exercise, first suppose that C is such that there is an edge disjoint path from every input fiber to every wavelength interchanger and an edge disjoint path from every wavelength interchanger to every output fiber. Suppose C contains $2k-2$ wavelength interchangers. For the purposes of this example, consider a demand set with $k-1$ demands with input wavelength λ_1 and output wavelength λ_1 and $k-1$ demands with input wavelength λ_2 and output wavelength λ_2. A valid routing of these demands would use $k-1$ wavelength interchangers to route the demands with input wavelength λ_1 and output wavelength λ_1. Furthermore it *could* route the $k-1$ demands with input wavelength λ_2 and output wavelength λ_2 through a disjoint set of $k-1$ wavelength interchangers. By the definition of strictly non-blocking, it must be possible to route another demand through C even if this is the current state of C. However, a new demand with input wavelength λ_1 and output wavelength λ_2 cannot be serviced by any of the wavelength interchangers in C. Thus C cannot be strictly non-blocking.

In general, we cannot assume that C has this structure. Therefore we will play the part of the adversary that is attempting to force C to reveal $2k-1$ wavelength interchangers. Initially, we do not know anything about the structure of the cross-connect. Each time we place a demand d on the cross-connect, the route r that C uses to satisfy d becomes known to us. Suppose that at some point there is a set of demands D on the cross-connect and a routing R of D. Furthermore, let d be a demand that is valid with respect to D. If there is a route r that could be used to satisfy d and r is known to us at that point, then we can also insist that the cross-connect use r to route d.

Therefore when proving our lower bounds we will sometimes insist that a demand be routed according to a particular route that is already known to exist.

3.2 Existence of Minimum Cost Split Cross-connects

In this section we prove that for the case of homogeneous $k \times k$ WDM cross-connects it is sufficient to consider only homogeneous $k \times k$ WDM *split* cross-connects. We reduce our problem to considering only strictly non-blocking $k \times k$ WDM split cross-connects, through the following steps.

1. We define the set *Long* of strictly non-blocking cross-connects that contain at least one path with multiple wavelength interchangers along the path.

2. We prove that *Long* cannot contain a strictly non-blocking 2×2 WDM cross-connect with 2 wavelength interchangers. This provides a base case for our inductive proof.

3. We then consider the smallest $k > 2$ such that there exists a strictly non-blocking $k \times k$ WDM cross-connect $C \in Long$ with fewer than $2k-1$ wavelength interchangers and show that C can be reduced to a smaller strictly non-blocking $k' \times k'$ WDM cross-connect C' with the property that C' has fewer than $2k'-1$ wavelength interchangers.

4. Thus, 2 and 3 together imply that if there exists a strictly non-blocking $k \times k$ WDM cross-connect $C \in Long$ with fewer than $2k-1$ wavelength interchangers, then we can reduce it to a strictly non-blocking $k' \times k'$ WDM cross-connect $C' \notin Long$ with fewer than $2k'-1$ wavelength interchangers.

More precisely, we define *Long* to 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 $w_P > 1$ wavelength interchangers to some output node $b \in O$. We say P is a *long* path. See Figure 3-1.

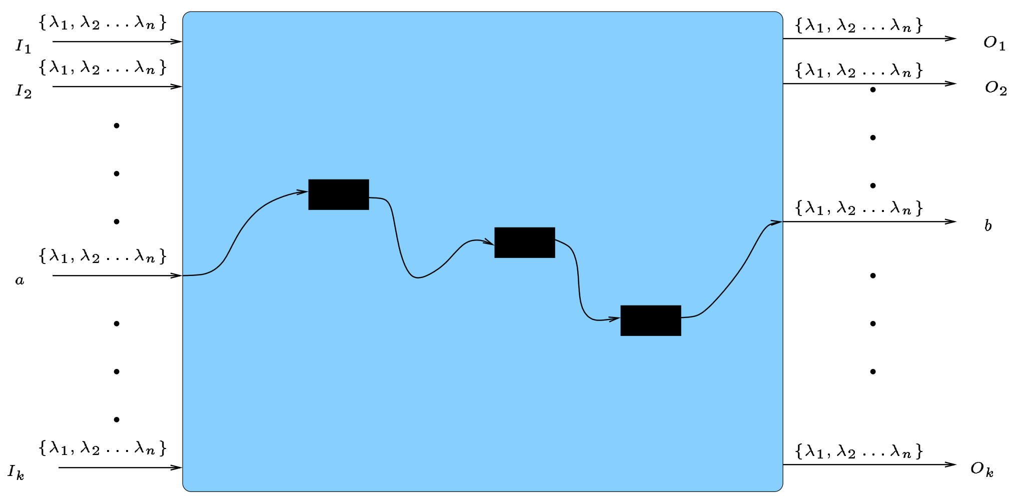

Figure 3-1: A long path in C with 3 wavelength interchangers.

To reduce a cross-connect $C \in Long$ to a smaller strictly non-blocking cross-connect, we make d demands from the input fiber a to the output fiber b. By assumption, there exists a long path P from a to b. Therefore we can insist that C route these n demands along P. As long as these demands are routed along P, no other demands on C can use any part of P. Therefore all future demands will be handled by other parts of C. Thus, C is then equivalent to a smaller cross-connect that has one less input fiber and one less output fiber and at least 2 fewer wavelength interchangers.

Formally, for any $C \in Long$ with n wavelengths supported on each input and output fiber and any long path $P \in C$ from input fiber a to output fiber b, we define the operation **Fill**(C, P, n) that routes demands $(a, \lambda_i, b, \lambda_i)$ for all $1 \leq i \leq n$ along P with constant wavelength assignment λ_i.

Lemma 3 *For any $C \in Long$, the result of the operation **Fill**(C, P, n) is that no new demands can be routed through any wavelength interchanger on P.*

Proof: Let D be the set of demands made by the operation $\mathbf{Fill}(C, P, n)$. Since $\mathbf{Fill}(C, P, n)$ routes one demand per wavelength, with constant wavelength wavelength assignment, along P, all wavelengths on all fibers in P are used for demands in D. Therefore there are no available supported wavelengths for any input or output fibers for any wavelength interchangers on P. ■

We now use Lemma 3 to show that any strictly non-blocking 2×2 WDM cross-connect $C \in Long$ must contain at least 3 wavelength interchangers. This provides a base case for our induction.

Lemma 4 *There does not exist a 2×2 WDM cross-connect $C \in Long$ with fewer than three wavelength interchangers.*

Proof: Let $n \geq 2$ be the number of wavelengths supported on each input and output fiber. Since $C \in Long$ it must have a path P with $w > 1$ wavelength interchangers. Therefore C must have exactly 2 wavelength interchangers and they must be on P. We start by performing $\mathbf{Fill}(C, P, n)$. Since we have only made demands from one of the input fibers to one of the output fibers, there is still an input fiber, say $e \in I$, and an output fiber, say $f \in O$ with all n wavelengths free. Consider a new demand $d = (e, \lambda_i, f, \lambda_j)$ where $1 \leq i, j \leq n$ and $i \neq j$. Since both e and f were previously not used in any demands, d is a valid demand on C. Furthermore since $\lambda_i \neq \lambda_j$, C must route d through a wavelength interchanger. However, by Lemma 3, after performing $\mathbf{Fill}(C, P, n)$ there are no wavelength interchangers with λ_i free on the input fiber and λ_j available on the output fiber. Thus C does not contain a route for d and therefore C is not strictly non-blocking. ■

Thus we now focus our attention on cross-connects in *Long* with $k > 2$. We show that we can reduce a strictly non-blocking $k \times k$ WDM cross-connect $C \in Long$ with $m < 2k - 1$ wavelength interchangers to a strictly non-blocking $k' \times k'$ WDM cross-connect with $m' < 2k' - 1$ wavelength interchangers.

Theorem 5 *If for some $k > 2$, there exists a $k \times k$ WDM cross-connect $C \in Long$ that has fewer than $2k - 1$ wavelength interchangers, then for some k', where $1 < k'$, there exists a strictly non-blocking $k' \times k'$ WDM cross-connect $C' \notin Long$ that has fewer than $2k' - 1$ wavelength interchangers.*

Proof: By Lemma 4, if there exists a $k \times k$ WDM cross-connect $C \in Long$ with fewer than $2k - 1$ wavelength interchangers, then $k > 2$. Let $k > 2$ be the smallest number such that there exists some $k \times k$ WDM cross-connect $C \in Long$ with $m < 2k - 1$ wavelength interchangers. Let w_P be the number of wavelength interchangers on a long path P from input fiber a to output fiber b that must exist in C by the definition of the class *Long* of cross-connects. Let n be the number of supported wavelengths.

We will use $\mathbf{Fill}(C, P, n)$ to route n demands along P. Once these demands have been routed along P, no new demands can use any portion of P. Thus any new demand on C must be routed through other sections of C. Therefore we show that C, with these n demands routed along P, is equivalent to a strictly non-blocking $k' \times k'$ WDM cross-connect C' with one less input fiber, one less output fiber and w_p fewer wavelength interchangers.

We start by performing $\mathbf{Fill}(C, P, n)$ and assume that the demands placed on C by $\mathbf{Fill}(C, P, n)$ are not removed. By Lemma 3 no new demands can be routed through wavelength interchangers on P. Thus consider the cross-connect C' obtained by the process

Modify(C, P) that is defined as follows. Remove input fiber a and output fiber 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'. See Figure 3-2 to see the change to C as a result of performing **Fill**(C, P, n) and **Modify**(C, P).

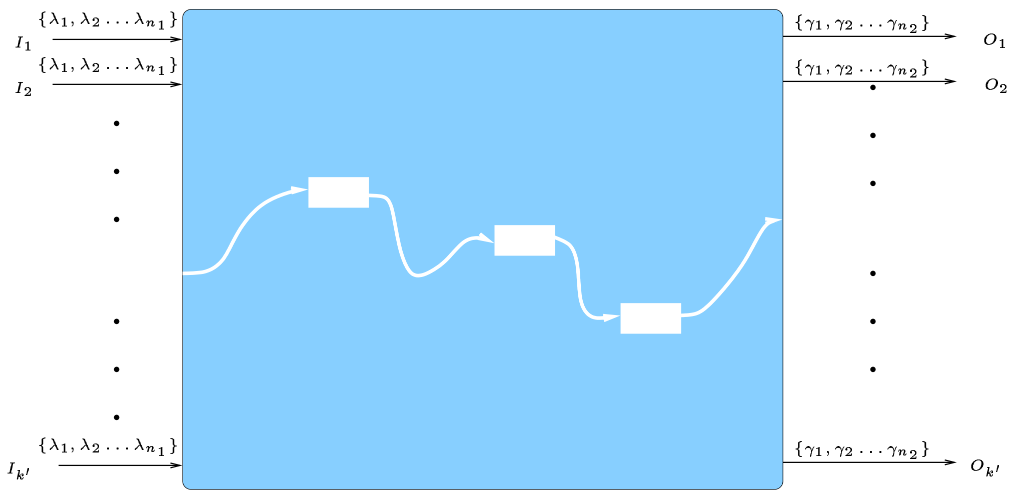

Figure 3-2: The result of performing **Fill**(C, P, n) and **Modify**(C, P).

Therefore since C is strictly non-blocking, C' must also be strictly non-blocking. Notice that since we have removed one input and output fiber from C to create C', C' is a $k' \times k'$ WDM cross-connect where $k' = k - 1$. Furthermore we removed at least two wavelength interchangers from C in **Modify**(C, P) and therefore the number of wavelength interchangers is $m - w_P < 2k' - 1$. Therefore C' is a $k' \times k'$ WDM cross-connect with fewer than $2k' - 1$ wavelength interchangers. Since k was assumed to be the smallest number such that there exists a $k \times k$ WDM cross-connect $C \in \mathit{Long}$ with fewer than $2k - 1$ wavelength interchangers and since $k' < k$ and $m' < 2k' - 1$, $C' \notin \mathit{Long}$. Thus $C' \notin \mathit{Long}$ is a strictly non-blocking $k' \times k'$ WDM cross-connect with fewer than $2k' - 1$ wavelength interchangers. ■

Theorem 5 says that if for some $k > 2$ there is an optimal strictly non-blocking $k \times k$ WDM cross-connect C with $m < 2k - 1$ wavelength interchangers and C is not a split cross-connect, then there is some $k' \geq 2$ for which there is a strictly non-blocking $k' \times k'$ WDM cross-connect C' with $m' < 2k' - 1$ wavelength interchangers. Furthermore, C' must be a split cross-connect. In the next section we will show that any strictly non-blocking $k \times k$ WDM split cross-connect must have at least $2k - 1$ wavelength interchangers.

3.3 Lower Bound for Split Cross-connects

In the last section we reduced the problem of showing that the optimal number of wavelength interchangers for any strictly non-blocking $k \times k$ WDM cross-connect is $2k - 1$ to showing that this is true for any strictly non-blocking $k \times k$ WDM split cross-connect. We now show that any strictly non-blocking $k \times k$ WDM split cross-connect must have $2k-1$ wavelength interchangers.

To do this we consider any strictly non-blocking WDM split cross-connect C with fewer than $2k-1$ 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 λ_2. Given this set of demands D and routing R we will then show that one additional valid demand $d \notin D$ exists with input wavelength λ_1 and output wavelength λ_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. Thus any strictly non-blocking $k \times k$ WDM split cross-connect must have at least $2k-1$ wavelength interchangers.

We say that a valid demand set D is *standard* if and only if

1. $|D| = 2k$ and

2. all of the demands in D use input wavelength λ_1 or λ_2 and output wavelength λ_1 or λ_2.

See Figure 3-3 for an example of a standard set of demands on C.

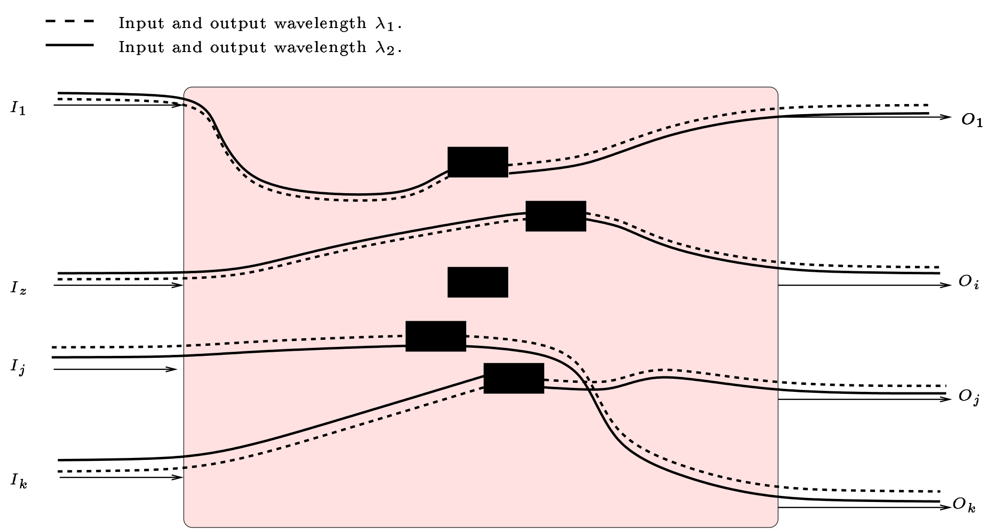

Figure 3-3: A set of standard demands D_0 on C.

Suppose we have a routing R_i for a standard set of demands D_i. Let W_i^B be the set of wavelength interchangers that service a demand with either input wavelength λ_1 or output wavelength λ_2. These wavelength interchangers are thought of as "blocking" demands with input wavelength λ_1 and output wavelength λ_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 λ_2.

If a routing R_i of a standard set of demands D_i uses fewer than $2k-1$ wavelength interchangers, then there must be at least two wavelength interchangers that are both in W_i^B and are each servicing two demands. As an example, assume WI_u and WI_v are two such wave-

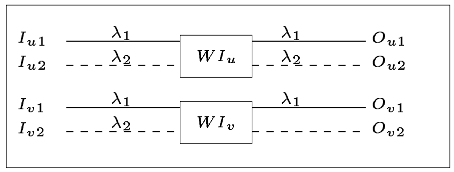

Figure 3-4: Two wavelength interchangers that each service two demands.

length interchangers. For this example, suppose $(I_{u1}, \lambda_1, O_{u1}, \lambda_1)$ and $(I_{u2}, \lambda_2, O_{u2}, \lambda_2)$ are the two demands that WI_u services. Furthermore, suppose $(I_{v1}, \lambda_1, O_{v1}, \lambda_1)$ and $(I_{v2}, \lambda_2, O_{v2}, \lambda_2)$ be the two demands that WI_v services. See figure 3-4.

Consider the following change to D_i.

- Remove $(I_{u1}, \lambda_1, O_{u1}, \lambda_1)$ and $(I_{v2}, \lambda_2, O_{v2}, \lambda_2)$ from D_i. See figure 3-5

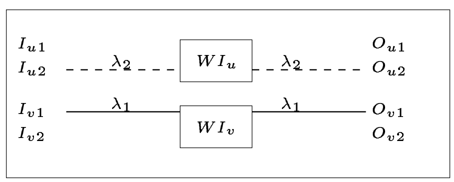

Figure 3-5: The effect of removing demands $(I_{u1}, \lambda_1, O_{u1}, \lambda_1)$ and $(I_{v2}, \lambda_2, O_{v2}, \lambda_2)$.

Notice that WI_u and WI_v both remain in W_i^B after we remove these two demands. However, I_{u1} now can ask for a new demand with input wavelength λ_1 and O_{v2} can now ask for a new demand with output wavelength λ_2. Now,

- add demand $(I_{u1}, \lambda_1, O_{v2}, \lambda_2)$ to D_i to create D_{i+1}. See figure 3-6

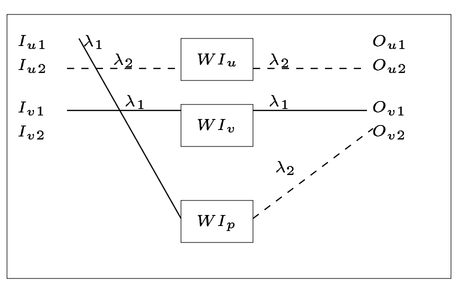

Figure 3-6: The effect of adding demand $(I_{u1}, \lambda_1, O_{v2}, \lambda_2)$.

This new demand is a valid demand with respect to D_i. However, since it has input wavelength λ_1 and output wavelength λ_2 it will not be able to be routed by any wavelength interchanger in W_i^B. Thus the route for this demand will force a new wavelength interchanger to service a demand with input wavelength λ_1 and output wavelength λ_2. The wavelength interchanger that services this demand will thus be added to W^B_{i+1}. By repeating this process we eventually will show that either C is not strictly non-blocking or there exists a set of standard demands D_t and a routing of these demands such that $|W^B_t| \geq 2k-1$ and thus C contains at least $2k-1$ wavelength interchangers.

We now formalize this construction. First we show that a standard set of demands D_0 exists and can be routed by C. Create k demands of the form $d_{i1} = (I_i, \lambda_1, O_i, \lambda_2)$ for $1 \leq i \leq k$. Clearly since C is strictly non-blocking, C must be able to satisfy these k demands. For $1 \leq i \leq k$, create a demand $d_{i2} = (I_i, \lambda_2, O_i, \lambda_1)$ and route along the same path as d_{i1}. Notice that there are $2k$ demands and each demand uses input wavelength λ_1 or λ_2 and output wavelength λ_1 or λ_2. Therefore this set of demands meets the definition of a standard set of demands.

Now we prove that any routing R_i of a standard set of demands D_i that uses $2k-g$ wavelength interchangers, will have g wavelength interchangers servicing 2 demands.

Lemma 6 *If C is a strictly non-blocking $k \times k$ WDM split cross-connect, D_i is a standard set of demands and R_i is a routing of the demands in D_i that uses $2k-g$ wavelength interchangers, where $g > 0$, then g wavelength interchangers in W^B_i will each service two demands both of whose input wavelengths are λ_1 and λ_2 and whose output wavelengths are λ_1 and λ_2.*

Proof: The $2k$ demands in D_i use only input wavelengths λ_1 and λ_2 and output wavelengths λ_1 and λ_2. As a result no wavelength interchanger can service more than two of these demands. This implies that g wavelength interchangers in W^B_i or W^F_i 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 λ_2 and therefore is by definition in W^B_i. ■

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 λ_2. Notice that this is equivalent to showing that if C has no more than $2k-g$ 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 $|W^B_i| = 2k-g$ and $|W^F_i| = 0$.

Let $WI_j \in W^B_i$ be a wavelength interchanger that services exactly two demands d_1 and d_2. Suppose these demands have the form $d_1 = (I_1, \lambda_1, O_2, \lambda_2)$ and $d_2 = (I_2, \lambda_2, O_1, \lambda_1)$. The operation **Uncross**(WI_j) is defined to have the effect of changing these two demands to be $(I_1, \lambda_1, O_1, \lambda_1)$ and $(I_2, \lambda_2, O_2, \lambda_2)$ and routing these new demands exactly as d_1 and d_2 were routed while keeping all other demands and routes unchanged. Recall that we can force the cross-connect to use a known route for a new demand in order to construct a particular routing. 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 $d_1 = (I_1, \lambda_1, O_2, \lambda_1)$ and $d_2 = (I_2, \lambda_2, O_1, \lambda_2)$, then **Uncross**$(WI_j)$ has no effect. See Figure 3-7.

Next we define the operation **Block**$(C, (D_i, R_i))$ where C is a strictly non-blocking $k \times k$ WDM split cross-connect, D_i is a standard set of demands and R_i is a routing of

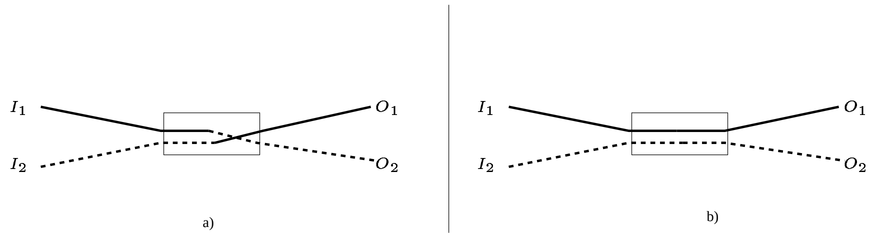

Figure 3-7: In *a* the wavelength interchanger is shown before the operation **Uncross** is performed. The result of **Uncross** is shown in *b*.

the demands in D_i that uses fewer than $2k-1$ 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 routing R_{i+1} for D_{i+1} such that $W_{i+1}^B > W_i^B$. In order to achieve this, **Block**$(C,(D_i,R_i))$ will use the demands serviced by two of the wavelength interchangers in W_i^B that must be servicing two demands by Lemma 6 to create a new demand with input wavelength λ_1 and output wavelength λ_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 6, each service two demands.
2. **Uncross**(WI_u) and **Uncross**(WI_v).
3. Let $(I_{u1},\lambda_1,O_{u1},\lambda_1)$ and $(I_{u2},\lambda_2,O_{u2},\lambda_2)$ be the two resulting demands that WI_u services. Let $(I_{v1},\lambda_1,O_{v1},\lambda_1)$ and $(I_{v2},\lambda_2,O_{v2},\lambda_2)$ be the two resulting demands that WI_v services.
4. Remove $(I_{u1},\lambda_1,O_{u1},\lambda_1)$ and $(I_{v2},\lambda_2,O_{v2},\lambda_2)$ from D_i and route all remaining demands according to R_i to create D_{i+1} and R_{i+1}.
5. Add $(I_{v2},\lambda_2,O_{u1},\lambda_1)$ to D_{i+1} and add the route chosen by C for this demand to R_{i+1}.
6. Add $(I_{u1},\lambda_1,O_{v2},\lambda_2)$ to D_{i+1} and add the route chosen by C for this demand to R_{i+1}.
7. Return (D_{i+1},R_{i+1}).

We start with D_0 and R_0 as defined above and then inductively define $(D_{i+1},R_{i+1}) =$ **Block**$(C,(D_i,R_i))$ as long as R_i uses fewer than $2k-1$ wavelength interchangers. Consider **Block**$(C,(D_i,R_i))$. By Lemma 6 and the assumption that R_i uses fewer than $2k-1$ 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 in Step 2. This does not change any of the wavelengths 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 λ_2. Thus WI_v and WI_u remain in W_i^B.

Since C is strictly non-blocking, a valid route for the new demand in Step 6 must exist. A wavelength interchanger in W_i^F must be used to service this demand since all wavelength

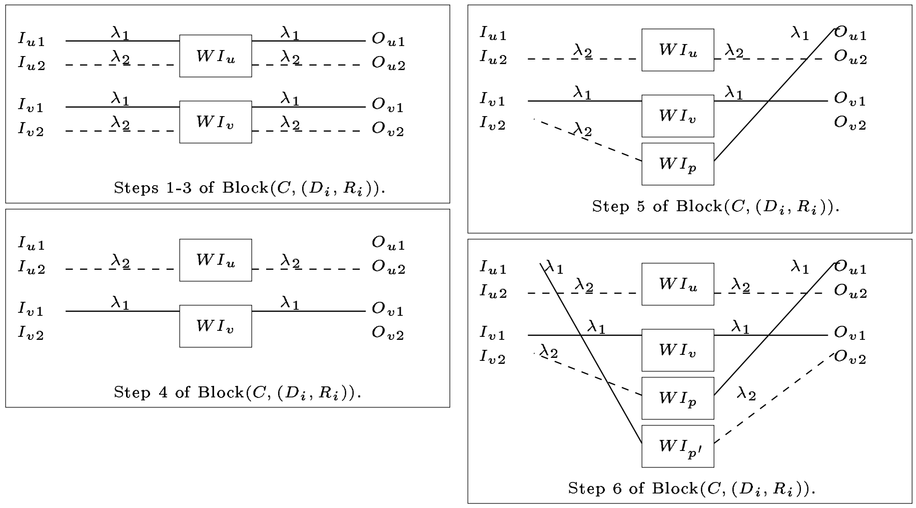

Figure 3-8: The steps in Block$(C, (D_i, R_i))$.

interchangers in W_i^B are servicing a demand with either input wavelength λ_1 or output wavelength λ_2. Since the demand requested in Step 6 has both input wavelength λ_1 and output wavelength λ_2, the wavelength interchanger that services this new demand will be in W_{i+1}^B. Notice that the demand in Step 5 insures that there are a total of $2k$ demands on C. Hence, after Step 6, D_{i+1} is a standard set of demands and R_{i+1} is a routing for those demands such that $|W_{i+1}^B| = |W_i^B| + 1$ and $|W_{i+1}^F| = |W_i^F| - 1$. Thus we can conclude:

Lemma 7 *If C is a strictly non-blocking $k \times k$ WDM split cross-connect, D_i is a standard set of demands and R_i is a routing of the demands in D_i that uses fewer than $2k-1$ wavelength interchangers, then* **Block**$(C, (D_i, R_i))$ *can be executed and* $(D_{i+1}, R_{i+1}) =$ **Block**$(C, (D_i, R_i))$ *where D_{i+1} is a standard set of demands and R_{i+1} is a routing of the demands in D_{i+1}. Also, D_{i+1} and R_{i+1} are such that $|W_{i+1}^B| = |W_i^B| + 1$ and $|W_{i+1}^F| = |W_i^F| - 1$.*

We now use Lemma 7 to show that there must be $2k-1$ wavelength interchangers in any strictly non-blocking $k \times k$ WDM split cross-connect.

Theorem 8 *For any strictly non-blocking $k \times k$ WDM split cross-connect there must be a standard set of demands D_t and a standard routing R_t of D_t that uses at least $2k-1$ wavelength interchangers.*

Proof: By contradiction. Let C be a strictly non-blocking $k \times k$ WDM split cross-connect. As noted earlier, for such a cross-connect there exists a standard set of demands D_0 with routing R_0. Let $M = 2k-1$. If R_0 uses at least M wavelength interchangers then we are done. Otherwise, by Lemma 7, we can repeatedly apply **Block** to produce a series of pairs $(D_{i+1}, R_{i+1}) =$ **Block**$(C, (D_i, R_i))$ where D_i is a standard set of demands and R_i is a 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 7 we know that $|W_{i+1}^F| = |W_i^F| - 1$ because the demand in Step 7 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 7 guarantees that we can perform **Block**$(C, (D_t, R_t))$. However, the demand in Step 7 must be serviced by a wavelength interchanger in W_t^F. Therefore Step 7 can not be completed in **Block**$(C, (D_t, R_t))$. Since C is strictly non-blocking it must be that R_t uses $2k - 1$ or more wavelength interchangers. ■

Theorem 8 says that any strictly non-blocking $k \times k$ WDM split cross-connect must have $2k - 1$ wavelength interchangers. Combining this with Theorem 5 we arrive at the following result for homogeneous cross-connects.

Corollary 9 *Any strictly non-blocking homogeneous $k \times k$ WDM cross-connect with $n > 1$ wavelengths must have at least $2k - 1$ wavelength interchangers.*

3.4 Other Demand Models for Homogeneous Cross-connects

In presenting the results of this chapter we have used a demand model in which the input and output wavelengths are specified as part of the demand. In this section we show that homogeneous cross-connects require $2k - 1$ wavelength interchangers even under less restrictive demand models. To prove this lower bound we must assume that $n \geq 4$. Although this is slightly weaker than our assumption in the previous chapters that $n \geq 2$, it covers all cases that arise in practice.

For a fiber a and a routing R of existing demands, we define $f(a, R)$ to be the set of wavelengths not in use on a by any route $r \in R$. We define a *general demand* $d = (I_x, O_y)$ to be a request for a connection from input fiber I_x to output fiber O_y. A *general route* r for d is a path from input fiber I_x to output fiber O_y such that the connection starts on an unused input wavelength, $\lambda_y \in f(I_x, R)$, stays on the same wavelength as long as it does not pass through a wavelength interchanger and ends on unused output wavelength $\lambda_w \in f(O_y, R)$. We say that a cross-connect uses a *general demand model* if it supports general demands and general routes.

Note that considering a strictly non-blocking cross-connect in a general demand model actually weakens the definition of a strictly non-blocking cross-connect. Since a strictly non-blocking cross-connect must route any new demand regardless of the set of current routes in use and the wavelength assignment of those routes, one can view a demand model that specifies the input and output wavelengths as part of the demand as representing a worse case adversary that requires not only a new connection from some input fiber to some output fiber but also chooses the *worst* wavelength assignment for the route chosen in the cross-connect. By allowing the cross-connect to choose the input and output wavelengths, the cross-connect is allowed to plan the choice of wavelengths. Therefore a cross-connect that supports only a general demand model is, to some extent, a hybrid between a strictly non-blocking cross-connect and a wide-sense non-blocking cross-connect. We consider such cross-connects and prove that allowing the cross-connect this added flexibility to choose the wavelength assignment of the routes does not change the number of wavelength interchangers needed in the cross-connect.

Note that although we allow the cross-connect to choose the input and output wavelengths, we still have the ability to force it to use a route that is known to satisfy a particular demand.

The outline for this section is similar to the earlier part of the chapter. We begin by reducing the problem to considering only split cross-connects. We then show that under certain conditions we can use a sequence of general demands to simulate a demand that specifies the input and output wavelengths as part of the demand. This allows us to prove that a strictly non-blocking $k \times k$ WDM split cross-connect supporting a general demand model requires $2k-1$ wavelength interchangers and thus any strictly non-blocking $k \times k$ WDM cross-connect supporting a general demand model requires $2k-1$ wavelength interchangers.

Thus, we begin this discussion showing that we can still reduce this problem to considering only split cross-connects. We assume throughout this section that $n \geq 4$. We define *LongGeneral* to be the set of strictly non-blocking WDM cross-connects using a general demand model that contain at least one directed path P from some input node $a \in I$ through $w_P > 1$ wavelength interchangers to some output node $b \in O$. We refer to P as a *long* path. See Figure 3-1. We define the operation **GeneralFill**(C, P, n) to route a set G of n demands from input fiber a to output fiber b along P.

Lemma 10 *For any long path P with $w_P \leq n-2$ and any routing chosen by the operation **GeneralFill**(C, P, n), there is a set S of demands such that*

1. *$S \subset G$,*
2. *S contains at most $n-1$ demands,*
3. *There is at least one wavelength λ_i that is used to route some demand in S on every section of fiber in P.*

Proof: By the assumptions of the lemma, P contains at most $n-2$ wavelength interchangers. Therefore any route r along P can change wavelength at most $n-2$ times. Therefore any particular demand can use at most $n-1$ wavelengths. This implies that for any wavelength λ_i there must be at least one demand $d \in G$ that is not assigned λ_i anywhere along its route. Let $S = G/\{d\}$. Since G contained n demands, each wavelength must be used by some demand on each section of P. Since d did not use λ_i on any section of P, the routes chosen for demands in S are such that for each section of fiber along P there is a route for some demand in S using λ_i. Therefore $S \subset G$ containing $n-1$ demands and the routes for demands in S are such that λ_i is used on a route for some demand in S along every section of fiber in P. ■

Using Lemma 10, we can now prove that there does not exist a *2* × *2* WDM cross-connect with at most 2 wavelength interchangers such that those wavelength interchangers are on a long path. We will use this as a base case for a reduction that shows that we can always consider split cross-connects when showing that any $k \times k$ WDM cross-connect must have $2k-1$ wavelength interchangers. Since a *1* × *1* WDM cross-connect cannot both have less than $2k-1=1$ wavelength interchanger *and* a long path, we must use a *2* × *2* WDM cross-connect as our base case for the reduction.

Theorem 11 *There does not exist a 2 × 2 WDM cross-connect $C \in$ LongGeneral with fewer than three wavelength interchangers.*

Proof: Assume that there exists a $C \in \mathit{LongGeneral}$ with 2 wavelength interchangers and $n \geq 4$. Let P be a long path from input fiber a to output fiber b with both wavelength interchangers in C on P. First perform **GeneralFill**(C, P, n) to route a set G of n demands from a to b. Notice that after this has been done all wavelength interchangers along P have every wavelength in use on their input and output fibers. Since all the wavelength interchangers in C are on P, any new demands must not change wavelength. Let α be the other input fiber and β be the other output fiber. Suppose we now ask for n demands from α to β. Either the routing for these demands must use constant wavelength assignments for these demands or it contradicts the assumption that C has exactly 2 wavelength interchangers. Therefore assume that these new demands are routed with constant wavelength assignment.

By Lemma 10 there exists $S \subset G$ such that S contains at most $n - 1$ demands and one demand in S is routed on wavelength λ_i on each section of fiber along P. Remove the demand from α to β that is routed with constant wavelength assignment λ_i. See Figure 3-9

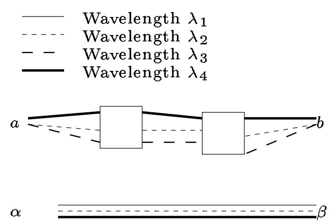

Figure 3-9: A 2×2 WDM cross-connect C with $n = 4$ and a path P from input fiber a to output fiber b such that both wavelength interchangers in C are on P. There are currently 3 demand routed along P such that each section of P has at least one demand routed on wavelength λ_3. There are 3 demands between α and β such that there is no demand from α to β routed along wavelength λ_3.

Furthermore, remove all demands along P that are not in S. At least one demand must be removed from P and that demand must have been routed such that it didn't use λ_i along any fiber in P. In particular it cannot have used λ_i on output fiber b. Now add a demand from α to b. Since λ_i is the only wavelength available on α, any route for this demand must start with wavelength assignment λ_i. Since λ_i is not available on output fiber b this demand must change wavelength. However, each fiber along P and hence each input fiber to each wavelength interchanger along P already has a demand in S routed on it that is using wavelength λ_i. Therefore there is no wavelength interchanger that can be used to service this new demand and as a result either $C \notin \mathit{LongGeneral}$ or C must contain 3 or more wavelength interchangers which violates the assumption that C has at most 2 wavelength interchangers. ■

Recall that Theorem 5 considered the smallest k for which there was a $k \times k$ WDM cross-connect $C \in \mathit{Long}$ that had fewer than $2k - 1$ wavelength interchangers. It then used **Fill**(C, P, n) to fill the path P with n demands. Once this path was filled with these demands, C was functionally equivalent to a $k' \times k'$ WDM cross-connect $C' \notin \mathit{Long}$ with fewer than $2k' - 1$ wavelength interchangers. Theorem 11 provides the base case for a similar reduction since it shows that there does not exist a cross-connect $C \in \mathit{LongGeneral}$ with $k = 2$ and at most 2 wavelength interchangers. Furthermore the effect of **GeneralFill**(C, P, n) when C is a cross-connect supporting a general demand model is to completely fill path P

such that no new demands can use any of the fibers or any wavelength interchanger along P. Thus the following is a straightforward extension of Theorem 5.

Corollary 12 *If for some $k > 2$, there exists a $k \times k$ WDM cross-connect $C \in$ LongGeneral with $n \geq 4$ that has fewer than $2k-1$ wavelength interchangers, then for some k', where $1 < k' < k$, there exists a strictly non-blocking $k' \times k'$ WDM cross-connect $C' \notin$ LongGeneral that has fewer than $2k'-1$ wavelength interchangers.*

Corollary 12 says that we need only consider whether there exists a strictly non-blocking $k \times k$ WDM *split* cross-connect supporting a general demand model that uses fewer than $2k-1$ wavelength interchangers. Therefore recall the version of Block$(C, (D_i, R_i))$ that requires that demands specify the input and output wavelengths. We assume that R_i uses at most $2k-2$ wavelength interchangers and every demand in D_i has input wavelength λ_1 or λ_2 and output wavelength λ_1 or λ_2.

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

1. Take two wavelength interchangers, WI_u and WI_v, in W_i^B that, by Lemma 6, each service two demands.
2. **Uncross**(WI_u) and **Uncross**(WI_v).
3. Let $(I_{u1}, \lambda_1, O_{u1}, \lambda_1)$ and $(I_{u2}, \lambda_2, O_{u2}, \lambda_2)$ be the two resulting demands that WI_u services. Let $(I_{v1}, \lambda_1, O_{v1}, \lambda_1)$ and $(I_{v2}, \lambda_2, O_{v2}, \lambda_2)$ be the two resulting demands that WI_v services.
4. Remove $(I_{u1}, \lambda_1, O_{u1}, \lambda_1)$ and $(I_{v2}, \lambda_2, O_{v2}, \lambda_2)$ from D_i and route all remaining demands according to R_i to create D_{i+1} and R_{i+1}.
5. Add $(I_{v2}, \lambda_2, O_{u1}, \lambda_1)$ to D_{i+1} and add a valid route for this demand to R_{i+1}.
6. Add $(I_{u1}, \lambda_1, O_{v2}, \lambda_2)$ to D_{i+1} and add a valid route for this demand to R_{i+1}.
7. Return (D_{i+1}, R_{i+1}).

Consider Step 5 which asks for a route from I_{v2} to O_{u1} such that the demand is routed on input wavelength λ_2 on input fiber I_{v2} and output wavelength λ_1 on output fiber O_{u1}. In the current model we cannot specify the input or output wavelength. Therefore we need to replace Step 5 with a sequence of general demands $\{d_1, d_2, \ldots d_j\}$ such that the final general demand, d_j, is from input fiber I_{v2} to output fiber O_{u1} and the routes for the demands $\{d_1, d_2 \ldots d_{j-1}\}$ require that demand d_j uses input wavelength λ_2 and output wavelength λ_1. Once d_j has been routed through C, we can remove demands $\{d_1, d_2 \ldots d_{j-1}\}$ and be left with the same result as Step 5. In order to show that this is possible, we will use the fact that the set of demands D_i and routing R_i of D_i are such that all demands are routed using either input wavelength λ_1 or λ_2 and either output wavelength λ_1 or λ_2.

More generally, for a set of demands D we will define a routing R of D to be **Constrained**(λ_i, λ_j) if for every demand $d \in D$, the route r for d uses either λ_i or λ_j as the input wavelength and either λ_i or λ_j as the output wavelength. Notice that the advantage of maintaining a **Constrained**(λ_i, λ_j) routing is that we are guaranteed that all other wavelengths are unused on all the fibers. Since $n \geq 4$ this implies that at most half of the wavelengths are in use on any given input or output fiber under any **Constrained**(λ_i, λ_j) routing.

Consider a cross-connect C and suppose there is a set of demands D where the routing R of D is **Constrained**(λ_i, λ_j). Furthermore assume that there is an input fiber I_a with $\lambda_i \in f(I_a, R)$ and there is an output fiber O_b with $\lambda_j \in f(O_b, R)$. We create the following operation **Force**$(C, (D, R), (I_a, \lambda_i, O_b, \lambda_j))$ that forces the route for a new general demand from I_a to O_b to use input wavelength λ_i and output wavelength λ_j. In describing **Force**$(C, (D, R), (I_a, \lambda_i, O_b, \lambda_j))$ we assume that $|f(I_a, R)| \geq |f(O_b, R)|$. The case when $|f(I_a, R) < |F(O_b, R)|$ is analogous.
Force$(C, (D, R), (I_a, \lambda_i, O_b, \lambda_j))$

1. Assume $|f(I_a, R)| \geq |f(O_b, R)|$.
2. Create $|f(O_b, R)|$ demands from input fiber I_a to output fiber O_b.
3. If a demand from Step 2 is routed on output wavelength λ_j, then remove it.
4. Let D' and R' be the current set of demands and current routing of these demands.
5. Create $|f(I_a, R')|$ demands from input fiber I_a to output fibers other than O_b.
6. Remove the demand that uses input wavelength λ_i on input fiber I_a and redefine D' and R' to be the current set of demands and routing of these demands.
7. Add $|f(O_b, R')|$ demands from input fibers other than I_a to output fiber O_b.
8. Remove the demand from output fiber O_b that is routed on output wavelength λ_j.
9. Add a demand from I_a to O_b to D' and add the route r for this demand to R'.
10. Remove any existing demands that were made in Steps 2, 5 and 7.
11. Return the current set of demands, D', and the current routing R' of D'.

Lemma 13 *Let C be a homogeneous strictly non-blocking $k \times k$ WDM split cross-connect with $n \geq 4$ available wavelengths. Let D be a set of demands and R a set of routes that are **Constrained**(λ_i, λ_j). If $\lambda_i \in f(I_a, R)$ and $\lambda_j \in f(O_b, R)$, then*

1. *$(D', R') =$ **Force**$(C, (D, R), (I_a, \lambda_i, O_b, \lambda_j))$ can be performed,*
2. *$D' = \{d \cup D\}$ where d is a demand from input fiber I_a to output fiber O_b,*
3. *$R' = \{r \cup R\}$ where r is the route for d using input wavelength λ_i and output wavelength λ_j,*
4. *and R' is **Constrained**(λ_i, λ_j).*

Proof: Notice that the case when $|f(I_a, R)| < |f(O_b, R)|$ is analogous to the case when $|f(I_a, R)| > |f(O_b, R)|$. Therefore without loss of generality we consider $|f(I_a, R)| \geq |f(O_b, R)|$. Note that since R was **Constrained**(λ_i, λ_j), each fiber can have at most 2 demands and therefore $n - 2 \leq |f(I_a, R)| \leq n$ and $n - 2 \leq |f(I_a, R)|$.

In Step 2 we make $|f(O_b, R)|$ demands. If $|f(I_a, R)| = |f(O_b, R)|$ then all free input wavelengths will be used on I_a. Otherwise, if $|f(I_a, R)| > |f(O_b, R)|$ then one wavelength on input fiber I_a will be unused. We then remove the demand that used output wavelength λ_j. Therefore before Step 5, $|f(I_a, R')| \leq 2$. Since R was **Constrained**(λ_i, λ_j) and since we have only added demands from output fiber O_b, all other output fibers must still have at least 2 wavelengths unused prior to Step 5. Therefore there are enough output wavelengths to make the requested demands in Step 5. After Step 6 there will be only one free wavelength, λ_i, on input fiber I_a and at most two free wavelengths on output fiber O_b. Therefore, by an analogous argument, there are enough free input wavelengths to make the demands requested in Step 7. Furthermore, prior to Step 7 we removed any demand from output fiber O_b that was routed on output wavelength λ_j. Therefore the demand we remove in Step 8 that is routed on output wavelength λ_j on output fiber O_b is guaranteed to use an input fiber other than I_a. Hence after step 8, input fiber I_a has only input wavelength λ_i available and output fiber O_b has only output wavelength λ_j available. Therefore the route for the demand made in Step 9 will use input wavelength λ_i and output wavelength λ_j. Finally, since we remove all demands added in Steps 2, 5 and 7, $D' = \{d \cup D\}$ where d is a demand from input fiber I_a to output fiber O_b and R' is **Constrained**(λ_i, λ_j). ■

Notice that Lemma 13 allows us to create a set of standard demands D and a set of routes R for those demands. Furthermore Lemma 13 allows us to replace the demands made in Steps 2, 5, and 6 of Block$(C, (D_i, R_i))$ with the appropriate sequence of general demands. Therefore we arrive at the following corollary to Theorem 8.

Corollary 14 *For any homogeneous strictly non-blocking $k \times k$ WDM split cross-connect using a general demand model with $k \geq 2$ and $n \geq 4$ there must be a standard set of demands D_t and a standard routing R_t of D_t that uses at least $2k-1$ wavelength interchangers.*

Combining the results in Corollaries 12 and 14 imply the following result for homogeneous cross-connects.

Corollary 15 *Any homogeneous strictly non-blocking $k \times k$ WDM cross-connect with $k \geq 2$ and $n \geq 4$ wavelengths supporting any demand model must have at least $2k-1$ wavelength interchangers.*

Chapter 4

Determining the Optimal Number of Wavelength Interchangers for Heterogeneous Split Cross-connects

In this chapter we focus on heterogeneous split cross-connects. Recall that a heterogeneous $k_1 \times k_2$ WDM cross-connect has k_1 input fibers that each support n_1 input wavelengths and k_2 output fibers that each support n_2 output wavelengths. Note that the set of input and output wavelengths can be completely disjoint. Therefore we now refer to the output wavelengths as $\{\gamma_1, \gamma_2 \ldots \gamma_{n_2}\}$. Without loss of generality, we assume that $k_1 \leq k_2$.

In Chapter 2 we presented a family of strictly non-blocking $k_1 \times k_2$ WDM cross-connects that use $\min(k_1 + k_2 - 1, n_1k_1)$ wavelength interchangers. In chapter 3 we showed that this is the optimal number of wavelength interchangers for any $k \times k$ WDM homogeneous cross-connect. In this chapter we extend the techniques of Chapter 3 to $k_1 \times k_2$ WDM heterogeneous split cross-connects. To show that any heterogeneous $k_1 \times k_2$ WDM split cross-connect requires $\min(k_1 + k_2 - 1, n_1k_1)$ wavelength interchangers we consider any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $m < \min(k_1 + k_2 - 1, n_1k_1)$ wavelength interchangers and show that there is a sequence of demands and a routing of those demands such that there is a final valid demand d that cannot be routed by the cross-connect. For the heterogeneous case we consider only split cross-connectsand only a demand model where the input and output wavelengths are specified as part of the demand.

The proof that any heterogeneous strictly non-blocking $k_1 \times k_2$ WDM split cross-connect must have $\min(k_1 + k_2 - 1, n_1k_1)$ wavelength interchangers has a similar flavor to the proof, presented in the previous chapter, that showed that any homogeneous $k \times k$ WDM cross-connect requires $2k - 1$ wavelength interchangers. We present the lower bound in two cases. First we consider $k_1 \times k_2$ WDM cross-connects where $n_1k_1 \leq k_2$. In this case we can create a set of demands that all use the same output wavelength and therefore must all use their own wavelength interchanger. Hence we arrive at a simple lower bound that says that any strictly non-blocking $k_1 \times k_2$ WDM cross-connect must have at least n_1k_1 wavelength interchangers if $n_1k_1 \leq k_2$. We then consider cross-connects where $n_1k_1 > k_2$. Our basic idea is to mimic the proof from Chapter 3 that showed that any $k \times k$ WDM cross-connect requires $2k-1$ wavelength interchangers. To achieve this for heterogeneous cross-connects we use two input wavelengths, $\{\lambda_1, \lambda_2\}$, two output wavelengths $\{\gamma_1, \gamma_2\}$ and a set of input and output fibers of equal size, say z, to force $2z-1$ wavelength interchangers to service either a demand with input wavelength λ_1 or a demand with output wavelength γ_2. We use the remaining $k_2 - z$ output fibers and the unused wavelengths on the input fibers to "hold" $k_2 - z$ wavelength interchangers. By making z as large as possible while still having enough available input wavelengths to make the $k_2 - z$ demands, we are able to show that any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect requires $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers.

4.1 Lower Bound for Heterogeneous Split Cross-connects

We approach the lower bound of $k_1 \times k_2$ WDM cross-connect as follows.

1. We show that if $k_2 \geq n_1 k_1$ then $n_1 k_1$ wavelength interchangers are necessary.
2. We show that when $k_2 < n_1 k_1$ any strictly non-blocking $k_1 \times k_2$ WDM cross-connect must have at least $\min(k_1 + k_2 - 1, n_1 k_1 - 1)$ wavelength interchangers .
3. We then consider the case when $(n_1 - 1)k_1 < k_2 < n_1 k_1$, where 2 only proved that $n_1 k_1 - 1$ wavelength interchangers are necessary, and show that in fact $n_1 k_1$ wavelength interchangers are necessary for any strictly non-blocking $k_1 \times k_2$ WDM cross-connect.

We begin this section by considering the case when $k_2 \geq n_1 k_1$ where we saw in Chapter 2 that there exists a strictly non-blocking WDM cross-connect that uses $n_1 k_1$ wavelength interchangers.

Theorem 16 *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 argument in the proof of Theorem 16 that any strictly non-blocking WDM cross-connect must have at least k_2 wavelength interchangers. In order to show that any strictly non-blocking WDM split cross-connect actually requires $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers, we first show that $\min(k_1, n_1 k_1 - k_2) + k_2 - 1 = \min(k_1 + k_2 - 1, n_1 k_1 - 1)$ wavelength interchangers are needed. To do this we consider any strictly non-blocking WDM split cross-connect C with fewer than $\min(k_1, n_1 k_1 - k_2) + k_2 - 1$ wavelength interchangers. For such a cross-connect we will create a demand set D and a routing R of D such that C must use each wavelength interchanger to service a demand with either input wavelength λ_1 or output wavelength γ_2. Given D and R we will then show that one additional valid demand $d \notin D$ exists with input wavelength λ_1 and output wavelength γ_2. By construction 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. Therefore the cross-connect will not have a wavelength interchanger to service the new demand. This implies that if a $k_1 \times k_2$ WDM split cross-connect is strictly non-blocking then it must have at least $\min(k_1 + k_2 - 1, n_1k_1 - 1)$ wavelength interchangers. For cases where this construction shows that $n_1k_1 - 1$ wavelength interchangers are necessary, we then start with the routing that uses $n_1k_1 - 1$ wavelength interchangers and then manipulate the demands to require n_1k_1 wavelength interchangers for any $k_1 \times k_2$ WDM cross-connect.

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 $\mathcal{O}^H$.

See Figure 4-1 for an example of a standard set of demands and a standard routing of those demands on a heterogeneous cross-connect C.

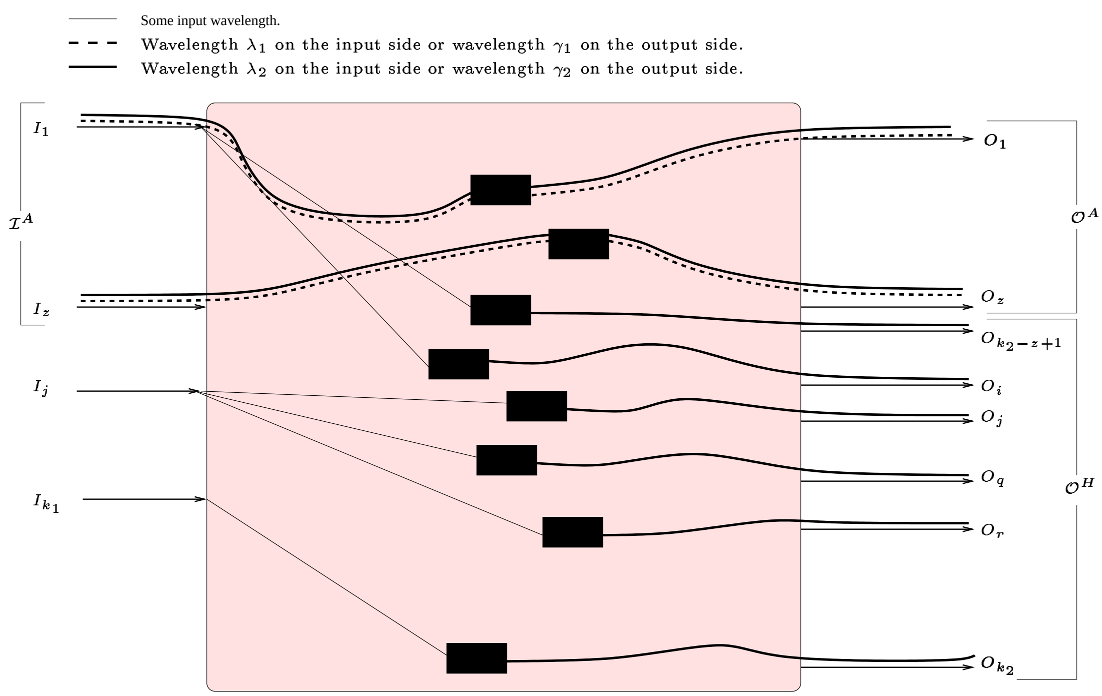

Figure 4-1: A set of standard demands, D_0, on C.

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.

As with our construction for homogeneous cross-connects we will assume that we have a routing R_i for a standard set of demands D_i. We will then successively transform the 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 for heterogeneous cross-connects we must 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 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 so the cross-connect will be unable to route other demands with output wavelength γ_2 through them. We modify the definition of W_i^B from the homogeneous case to 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. Finally, let W_i^F be the set of all other wavelength interchangers. This set of wavelength interchangers is 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 $k_2 < n_1 k_1$. 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 $1 \leq i \leq z$. 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 $d_{i2} = (a_i, \lambda_2, b_i, \gamma_2)$ routed along the same path as d_{i1}.

Now with these $2z$ demands on the cross-connect, we would like to create $k_2 - z$ demands such that each output fiber $h_j \in \mathcal{O}^H$ has a demand with output wavelength γ_2. First consider the case where $k_1 \leq n_1 k_1 - k_2$ and hence $z = k_1$. In this case $(n_1 - 1)k_1 \geq k_2$. Thus, $(n_1 - 2)k_1 \geq k_2 - k_1$. Since there are $n_1 k_1 - 2k_1$ input wavelengths unused on the set of input fibers, clearly there are enough input wavelengths available to make $k_2 - k_1$ new valid demands all using output wavelength γ_2. On the other hand if $z = n_1 k_1 - k_2$, then $n_1 k_1 - 2z = k_2 - z$ and therefore there are again enough available wavelengths on the input side to make $k_2 - z$ demands with output wavelength γ_2. Therefore we add $k_2 - z$ valid demands such that each fiber $h_j \in \mathcal{O}^H$ has a demand with output wavelength γ_2.

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 17 *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. Thus 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 a manipulation, similar to Block$(C, (D_i, R_i))$, that 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 $d_2 = (a_2, \lambda_2, b_1, \gamma_1)$. Recall the operation **Uncross**(WI_j) which 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. Again, 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. See Figure 3-7.

We define the operation **HeterogeneousBlock**$(C, (D_i, R_i))$ which is similar to **Block** $(C, (D_i, R_i))$ however now we assume that C is a heterogeneous 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. Each time we run **HeterogeneousBlock**$(C, (D_i, R_i))$ we 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 Chapter 3 we saw the operation **Block**$(C', (D_i, R_i))$ which performed a similar task for a homogeneous cross-connect C'. In **Block**$(C', (D_i, R_i))$ we did not have "extra" demands, D^H, that "held" the wavelength interchangers in W^H. For a heterogeneous cross-connect we use these demands to block the wavelength interchangers in W^H from servicing the two new demands that we create each iteration. In order to prevent either demand from being routed through one of the wavelength interchangers in W^H, we switch the output wavelength of the demands in D^H such that the demands in D^H always use the same output wavelength as the new demand that we are creating. See figure 4-2

After we define **HeterogeneousBlock**$(C, (D_i, R_i))$, we will show that it consists of a sequence of valid operations.

HeterogeneousBlock$(C, (D_i, R_i))$

1. Take two wavelength interchangers, WI_u and WI_v, in W_i^B that, by Lemma 17, 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}).

The initial set of demands on C is D_0 and they are routed according to R_0 as defined above. We inductively define $(D_{i+1}, R_{i+1}) =$ **HeterogeneousBlock**$(C, (D_i, R_i))$ as long as R_i uses fewer than $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ wavelength interchangers. Consider **HeterogeneousBlock**$(C, (D_i, R_i))$. We now prove that each step can be performed. Consider Steps 1 through 3.

1. Take two wavelength interchangers, WI_u and WI_v, in W_i^B that, by Lemma 17, 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.

By Lemma 17 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 wavelengths used along any of the fibers in C. In particular,

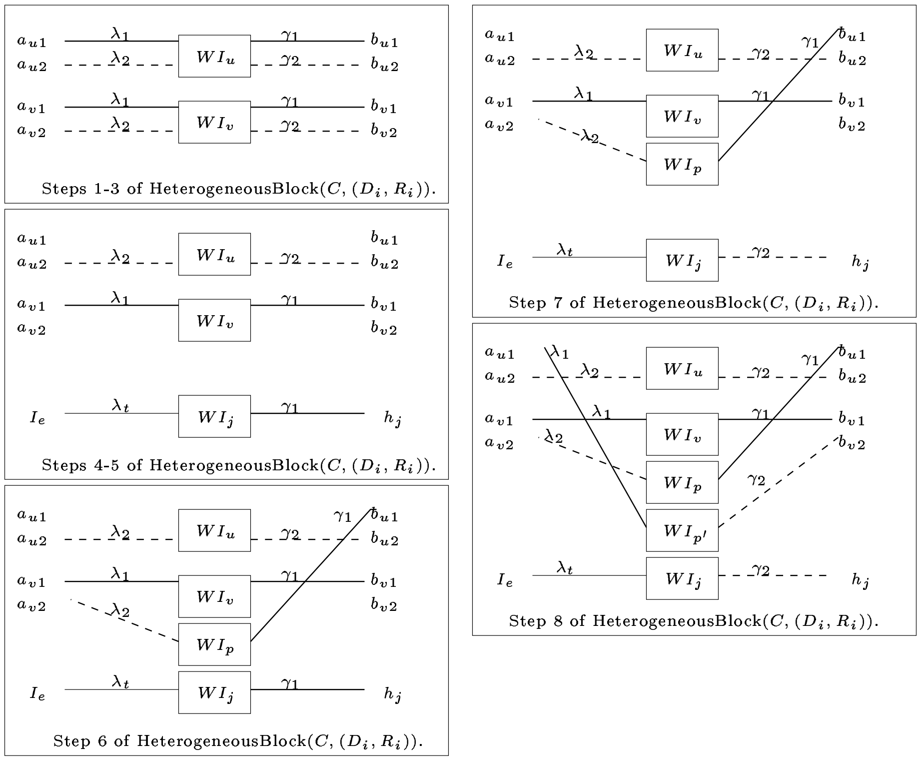

Figure 4-2: The steps in HeterogeneousBlock$(C, (D_i, R_i))$.

while this alters the set of demands, the new set of demands is still a standard set of demands.

In Step 4,

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}.

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. Thus WI_v and WI_u remain in W_i^B.

Next in Step 5,

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.

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.

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}.

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,

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.

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.

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}.

A valid route for the new demand in Step 8 must exist since C is strictly non-blocking. Furthermore 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. Thus, 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 $|W_{i+1}^F| = |W_i^F| - 1$. Hence we conclude:

Lemma 18 *If C is a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $k_2 < n_1 k_1$, 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*

1. **HeterogeneousBlock**$(C, (D_i, R_i))$ *can be executed,*
2. $(D_{i+1}, R_{i+1}) =$ **HeterogeneousBlock**$(C, (D_i, R_i))$,
3. D_{i+1} *is a standard set of demands,*
4. R_{i+1} *is a standard routing of the demands in* D_{i+1},
5. $|W^H|$ *remains fixed,*
6. $|W_{i+1}^B| = |W_i^B| + 1$ *and*
7. $|W_{i+1}^F| = |W_i^F| - 1$.

We now apply Lemma 18 to show that any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $k_1 \leq k_2 < n_1 k_1$ must have at least $\min(k_1, n_1 k_1 - k_2) + k_2 - 1$ wavelength interchangers.

Theorem 19 *For any strictly non-blocking $k_1 \times k_2$ WDM split cross-connect with $k_1 \leq k_2 < n_1 k_1$ there must be a standard set of demands D_t and a standard routing R_t of D_t that uses at least* $\min(k_1, n_1k_1 - k_2) + k_2 - 1 = \min(k_1 + k_2 - 1, n_1k_1 - 1)$ *wavelength interchangers.*

Proof: By contradiction. Let C be a strictly non-blocking $k_1 \times k_2$ WDM split cross-connect where $k_1 \leq k_2 < n_1k_1$. Start with the standard set of demands D_0 with standard routing R_0 which must exist. 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 18, we can repeatedly apply **HeterogeneousBlock** to produce a series of pairs $(D_{i+1}, R_{i+1}) =$ **HeterogeneousBlock**$(C, (D_i, R_i))$ where each D_i is a standard set of demands and each R_i is a standard routing of the demands in D_i. This sequence continues as long as R_i has the property that it uses fewer than M wavelength interchangers. By Lemma 18 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 18 guarantees that we can perform **HeterogeneousBlock**$(C, (D_t, R_t))$. 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 **HeterogeneousBlock**$(C, (D_t, R_t))$. Since C is strictly non-blocking it must be that R_t uses $\min(k_1, n_1k_1 - k_2) + k_2 - 1$ or more wavelength interchangers. ■

Theorem 19 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.

We could easily show that n_1k_1 wavelength interchangers are necessary, if we could continue executing $(D_{i+1}, R_{i+1}) =$ **HeterogeneousBlock**$(C, (D_i, R_i))$ until R_{i+1} uses $n_1k_1 - 1$ wavelength interchangers and $W_{i+1}^F = 0$. However, in general, it is possible that at some point, **HeterogeneousBlock**$(C, (D_i, R_i))$ 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 $W_{i+1}^F > 0$. Once this many wavelength interchangers in C are used to service a demand we can no longer run **HeterogeneousBlock**$(C, (D_{i+1}, R_{i+1}))$. Thus we need a procedure that is similar to **HeterogeneousBlock**$(C, (D_{i+1}, R_{i+1}))$ 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 17 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 20 *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 $z = n_1k_1 - k_2 < k_1$. The total number of demands is $k_2 + z = n_1k_1$. Therefore every input fiber has a demand for every input wavelength. Only z input fibers have demands with input wavelength λ_i in D^A. 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 assume 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 20 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 **HeterogeneousBlock2**$(C, (D_i, R_i))$ defined as follows.

HeterogeneousBlock2$(C, (D_i, R_i))$

1. Take one wavelength interchanger, WI_v, in W_i^B 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 on which d_u was routed.
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_i^F 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}).

Figure 4-3 shows the steps of HeterogeneousBlock2$(C, (D_i, R_i))$. In Lemma 21 we prove its correctness.

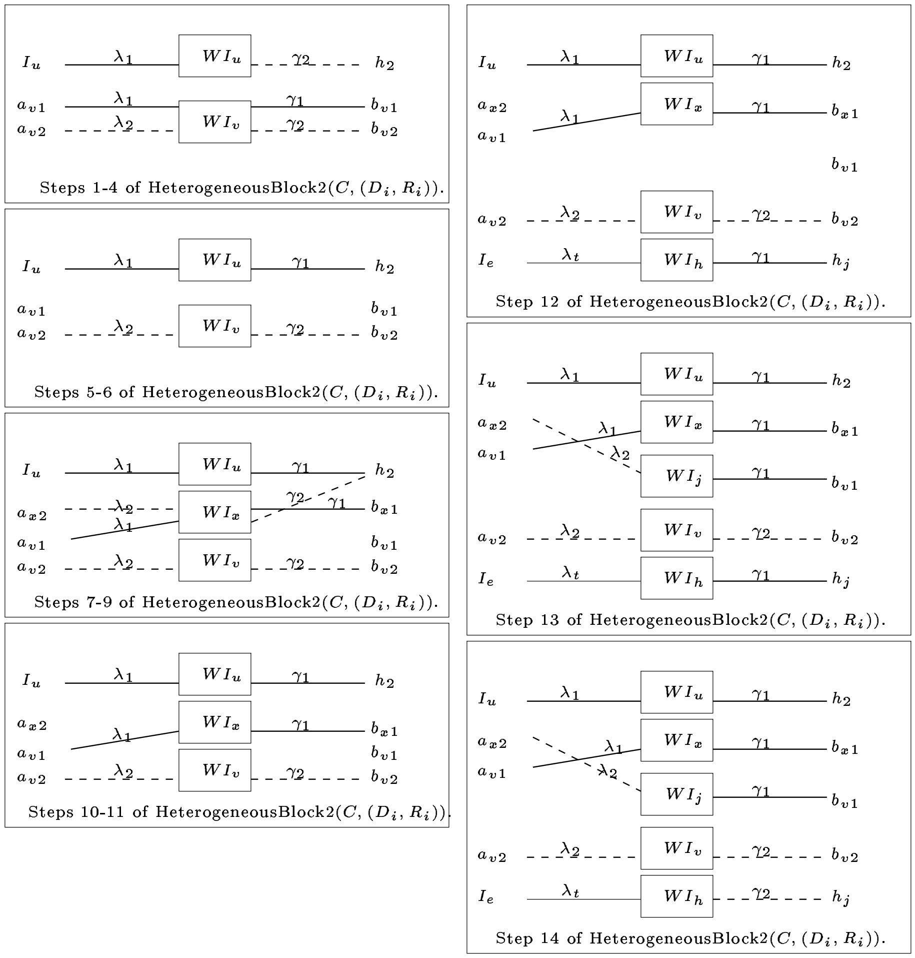

Figure 4-3: The steps in HeterogeneousBlock2$(C, (D_i, R_i))$.

Lemma 21 *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* **HeterogeneousBlock2**$(C, (D_i, R_i))$ *can be executed and* $(D_{i+1}, R_{i+1}) =$ **HeterogeneousBlock2**$(C, (D_i, R_i))$ *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, $|W_{i+1}^B| = |W_i^B| + 1$ and $|W_{i+1}^F| = |W_i^F| - 1$.*

Proof: We look at each step of **HeterogeneousBlock2**$(C, (D_i, R_i))$ and prove its correctness. Consider first Steps 1-4.

1. Take one wavelength interchanger, WI_v, in W_i^B 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.

See figure 4-4

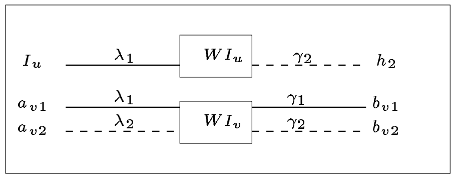

Figure 4-4: Steps 1-4 of HeterogeneousBlock2$(C, (D_i, R_i))$

Lemmas 17 and 20 show the existence of a wavelength interchanger $WI_v \in W_i^B$ 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. Thus consider Step 5.

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 on which d_u was routed.

After Step 5, when 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. After Step 6 input fiber a_{v1} will have input wavelength λ_1 free.

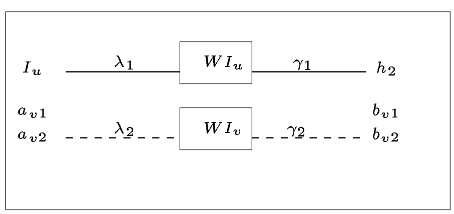

Figure 4-5: Steps 5 and 6 of HeterogeneousBlock2$(C, (D_i, R_i))$

6. Remove $d_{v1} = (a_{v1}, \lambda_1, b_{v1}, \gamma_1)$ and route all remaining demands according to R_i.

See figure 4-5.

We then add the demand,

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}.

from input fiber a_{v1} to output fiber h_2. See figure 4-6. This new demand uses input wavelength λ_1 and output wavelength γ_2. Therefore it must be serviced by a wavelength interchanger $WI_x \in W_i^F$ since all of the wavelength interchangers in W_i^B and W^H are

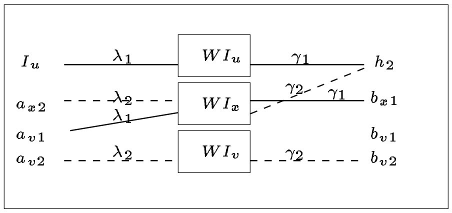

Figure 4-6: Steps 7-9 of HeterogeneousBlock2($C, (D_i, R_i)$)

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 every wavelength interchanger in C services a demand in D_i. Thus Steps 9 and 10 are,

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.

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 route demand d_q from Step 11,

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}.

along $r_q = p'_{v1}||p_{x1}$. See figure 4-7. Notice that this adds WI_x to W^B_{i+1}.

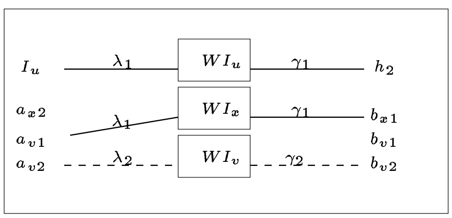

Figure 4-7: Steps 10 and 11 of HeterogeneousBlock2($C, (D_i, R_i)$)

Next in Step 12,

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.

we "switch" the output wavelength used by all demands in D^H. See figure 4-8. Since R_i is 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,

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}.

is not serviced by a wavelength interchanger in W^H. See figure 4-9.

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

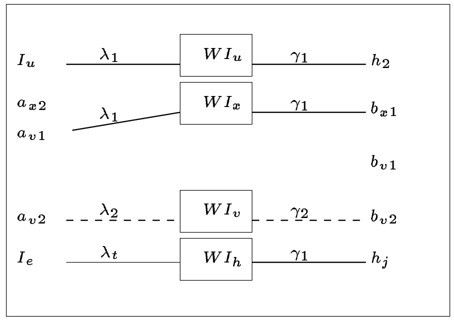

Figure 4-8: Step 12 of HeterogeneousBlock2$(C, (D_i, R_i))$

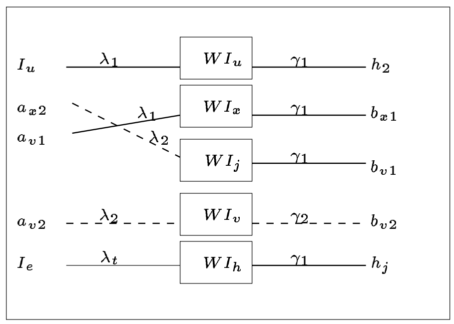

Figure 4-9: Step 13 of HeterogeneousBlock2$(C, (D_i, R_i))$

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, Step 14,

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.

switches 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. See figure 4-10.

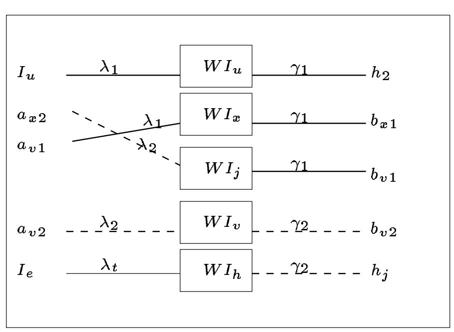

Figure 4-10: Step 14 of HeterogeneousBlock2$(C, (D_i, R_i))$

Notice that $|D_{i+1}| = k_2 + z$. Furthermore, $2z$ 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. The other $k_2 - z$ demands in D_{i+1} have output wavelength γ_2 and an output fiber in $\mathcal{O}^H$. Additionally 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 **HeterogeneousBlock2**$(C, (D_i, R_i))$ 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 18 we can repeatedly perform **HeterogeneousBlock**$(C, (D_i, R_i))$ until all $n_1k_1 - 1$ 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 21 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 **HeterogeneousBlock2**$(C, (D_t, R_t))$, 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 22 *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 16, 19 and 22 imply the following.

Theorem 23 *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.*

Chapter 5

Conclusions and Future Work

For a homogeneous $k \times k$ WDM cross-connect we have presented an optimal design using $2k-1$ wavelength interchangers. For heterogeneous cross-connects we have considered only the case of a $k_1 \times k_2$ WDM split cross-connect and have shown that $\min(k_1 + k_2 - 1, n_1 k_1)$ wavelength interchangers are necessary and sufficient. One obvious open question is to consider whether a non-split design is optimal for heterogeneous cross-connects. The reason we cannot extend Theorem 5 to heterogeneous cross-connects is that we need to be able to remove $n-1$ wavelength interchangers along a long path in some cases of heterogeneous cross-connects. However as with the homogeneous case we can only be sure to remove 2 wavelength interchangers. Therefore a new idea is required to prove that split cross-connects are optimal for heterogeneous cross-connects. Considering general demand models for heterogeneous cross-connects is also an interesting line of work. Another open area of work is to consider wide-sense non-blocking cross-connects. In [5] it is shown that wide-sense non-blocking $k \times k$ WDM split cross-connects with $n \geq O(k^2)$ available wavelengths require at least $2k-1$ wavelength interchangers. However, it is still open as to whether a split cross-connect is optimal for wide-sense non-blocking cross-connects. Furthermore for wide-sense non-blocking $k \times k$ WDM cross-connects with $2 < n < O(k^2)$ the optimal number of wavelength interchangers is unknown even if only split cross-connects are considered.

Various work has assumed that providing wavelength conversion capability for any demand passing through a node in a network would be independent of the degree of the node[15, 6]. However, it seems likely that the number of wavelength interchangers needed in a cross-connect placed at a node will still depend, to some extent on the degree of the node. Determining this relationship and the model that correctly captures the type of wavelength conversion capability necessary at a node in an arbitrary network was one motivation for this work and has been considered [11] but is largely still an open question.

Bibliography

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

[2] 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.

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

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

[5] P. Haxell, A. Rasala, G. Wilfong, and P. Winkler. Dynamic edge coloring with applications to optical cross-connect design. 2001.

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

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

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

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

[10] A. Rasala and G. Wilfong. Strictly non-blocking WDM cross-connects for heterogeneous networks. In *Proceedings of STOC*, 1999.

[11] A. Rasala and G. Wilfong. A new cost model for providing wavelength conversion capability. 2000.

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

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

[14] 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.

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