Wide-Sense Nonblocking WDM Cross-Connects

Penny Haxell[1⋆], April Rasala[2⋆⋆], Gordon Wilfong[3], and Peter Winkler[3]

[1] Department of Combinatorics and Optimization, University of Waterloo
Waterloo, Ont. Canada, N2L 3G1
pehaxell@math.uwaterloo.ca
[2] MIT Laboratory for Computer Science
Cambridge, MA, USA 02139
arasala@theory.lcs.mit.edu
[3] Bell Labs, Lucent Technologies
Murray Hill NJ, USA 07974
{gtw,pw}@research.bell-labs.com

Abstract. We consider the problem of minimizing the number of wavelength interchangers in the design of wide-sense nonblocking cross-connects for wavelength division multiplexed (WDM) optical networks. The problem is modeled as a graph theoretic problem that we call *dynamic edge coloring*. In dynamic edge coloring the nodes of a graph are fixed but edges appear and disappear, and must be colored at the time of appearance without assigning the same color to adjacent edges.
For wide-sense nonblocking WDM cross-connects with k input and k output fibers, it is straightforward to show that $2k-1$ wavelength interchangers are always sufficient. We show that there is a constant $c > 0$ such that if there are at least ck^2 wavelengths then $2k$–1 wavelength interchangers are also necessary. This improves previous exponential bounds. When there are only 2 or 3 wavelengths available, we show that far fewer than $2k-1$ wavelength interchangers are needed. However we also prove that for any $\varepsilon > 0$ and $k > 1/2\varepsilon$, if the number of wavelengths is at least $1/\varepsilon^2$ then $2(1-\varepsilon)k$ wavelength interchangers are needed.

1 Introduction

A wavelength division multiplexed (WDM) network employs multiple wavelengths in order to carry many channels in an optical fiber. A WDM network contains places at which multiple fibers come together. At these places, channels that have previously been routed along the same fiber may each need to be moved to different fibers and possibly also change wavelengths. Switching is, ideally, done by a WDM cross-connect that allows each incoming input channel to be routed to any (unused) output channel. To do this the cross-connect requires, among other things, expensive components called *wavelength interchangers* that permute the wavelengths on a fiber in any desired manner.

⋆ Partially supported by NSERC.
⋆⋆ Supported by a Lucent GRPW Fellowship.

R. Möhring and R. Raman (Eds.): ESA 2002, LNCS 2461, pp. 538–550, 2002.

Over time, the demands on the network change. Some connections are no longer needed and requests are made to add new connections. Any fiber in the network with an available wavelength can handle the addition of a new request by routing it on the unused wavelength. However, at cross-connects the interactions between requests can be more complicated. A cross-connect is said to be *wide-sense nonblocking* if there is an on-line algorithm that assures that it can always meet demands. (This is weaker than *strictly nonblocking* where the demands are never blocked even when previous demands have been routed arbitrarily).

Our goal here is to minimize the number of wavelength interchangers in the design of a wide-sense nonblocking cross-connect with k input fibers and k output fibers. It is easily seen that $2k-1$ wavelength interchangers suffice, even with greedy routing; we show that there is a constant $c > 0$ so that with ck^2 wavelengths, $2k-1$ wavelength interchangers are necessary as well, *regardless* of the routing algorithm. This improves previous exponential bounds.

On the positive side, in the case where there are only 2 or 3 wavelengths there is a significant reduction in the number of wavelength interchangers required. However, we also show that for any $\varepsilon > 0$ and $k > 1/2\varepsilon$, if there are at least $1/\varepsilon^2$ wavelengths then $2(1-\varepsilon)k$ wavelength interchangers are necessary.

This WDM cross-connect problem is shown to be equivalent to a dynamic edge coloring problem for bipartite multigraphs and the results are stated and derived in terms of this edge coloring problem.

2 Wavelength Division Multiplexing

In wavelength division multiplexing (WDM) an optical fiber or other medium carries many channels at once, subject to the constraint that each employs a different wavelength from some fixed set of Λ wavelengths. WDM systems greatly increase the available bandwidth of existing facilities, and are rapidly proliferating; systems with 80 wavelengths are becoming commonplace and systems with thousands are being contemplated.

Optimal use of bandwidth in a WDM network requires switches that can change the wavelength, as well as the fiber, on which a channel is carried [13, 14, 7, 16, 17]. A $k \times k$ *WDM cross-connect* should in theory be able to dynamically route up to Λk incoming channels on k fibers in any specified way onto k outgoing fibers, subject to the constraint that no two channels of the same wavelength are output on the same fiber. When the cross-connect is in operation, "demands" arrive and depart, and must be handled without knowledge of the future; each demand consists of an input channel (that is, an input fiber and a wavelength) and an output channel to which it must be linked.

The cost of a WDM cross-connect is dominated by the cost of the components, called "wavelength interchangers", that permute wavelengths on a fiber [23]. Thus our goal is to study WDM cross-connect designs that minimize the number of wavelength interchangers required to achieve certain nonblocking properties.

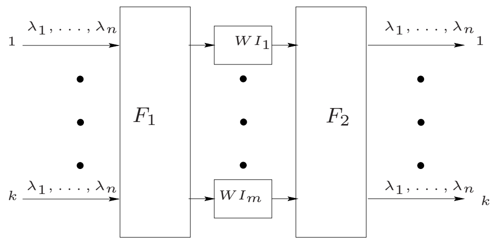

Fig. 1. A WDM split cross-connect with k input fibers and k output fibers

We consider an important class of WDM cross-connects known as *split cross-connects*, illustrated in Figure 1. In such cross-connects, any input channel can be routed to any wavelength interchanger not currently servicing a demand with the same input wavelength; and similarly any output channel can be routed from any wavelength interchanger not servicing a demand of the same output wavelength, regardless of how any previous demands have been routed [19, 20]. Thus in order to satisfy a demand the only decision necessary is to choose which available wavelength interchanger to use.

A demand for a connection from input channel I to output channel O is said to be *valid* if neither I nor O is part of an already routed demand. Demands for connections are requested and withdrawn over time. The nonblocking properties of a split WDM cross-connect are said to be *rearrangeably*, *wide-sense* or *strictly* nonblocking where

(i) "rearrangeably nonblocking" means that there exists an available wavelength interchanger to service any valid demand although the wavelength interchangers assigned to currently routed demands might have to be changed;
(ii) "wide-sense nonblocking" means that there exists an algorithm that always finds an available wavelength interchanger to service a valid demand assuming that all current assignments of wavelength interchangers to demands have been done using the same algorithm; and
(iii) "strictly nonblocking" means that there always exists an available wavelength interchanger to service any valid demand irrespective of how the previous assignments of wavelength interchangers to demands was performed.

The question of whether weakening the nonblocking constraint on traditional (i.e. non-WDM) cross-connects allows for less complex designs has been well studied. In traditional cross-connect design, the goal is to minimize the size of the cross-connect (i.e. the number of edges in the directed graph representing the connectivity in the cross-connect), and nonblocking properties are concerned with being able to route a valid demand avoiding edges used by previous demands. In this case, it has been shown that $\Omega(k \log k)$ is a lower bound on the size of a wide-sense nonblocking cross-connect [21] (in fact, it is actually shown to be a lower bound for the weaker rearrangeably nonblocking constraint). Also, it is known that $O(k \log k)$ is an upper bound for strictly nonblocking cross-connects [2] (and hence also an upper bound for wide-sense nonblocking cross-connects). That is, these bounds are tight (up to a constant factor) for both wide-sense and strictly nonblocking cross-connects and so there is no reduction in the size of a cross-connect to be gained by relaxing the nonblocking constraint to wide-sense. However, for more general kinds of demands (e.g. multicast demands, as in [3, 9]), there is a reduction in the required size of a wide-sense nonblocking cross-connect compared to a strictly nonblocking cross-connect.

Thus in the case of WDM cross-connects, it is again natural to study whether the weaker nonblocking properties allow for more efficient designs. It has been shown that for rearrangeably nonblocking split WDM cross-connects, k wavelength interchangers are necessary and sufficient [23]. The stronger property of strictly nonblocking was shown to have upper and lower bounds of $2k-1$ wavelength interchangers [19]. We consider the question of where the bounds for the intermediate case of wide-sense nonblocking lie. We begin by defining an equivalent graph edge coloring problem and then present our technical results in terms of the edge coloring problem.

3 The Graph-Theoretic Model

A *graph* here consists of a set of nodes and a multiset of edges; if there is at most one edge for each pair of nodes, we say that the graph is *simple*. In a *dynamic graph* edges appear and disappear over time (see Section 5 for the precise definition). An algorithm for edge coloring a dynamic graph must assign a color to each new edge presented without any knowledge of future additions or deletions. At all times adjacent edges must have different colors. The goal of the algorithm is to minimize the total number of colors ever assigned to edges of the graph.

Determining the number of necessary wavelength interchangers in a wide-sense nonblocking cross-connect can be cast as an edge coloring problem for dynamic bipartite graphs of fixed maximum degree. The set A of nodes on the left side of the graph represents the set of wavelengths available on each of the input fibers; the set B of nodes on the right side represents the set of wavelengths available on the output fibers. We define $n = \max(|A|, |B|)$ to be the *size* of the graph. (Normally the set of input wavelengths and the set of output wavelengths are the same, thus the bipartition is balanced.) An edge $\{u, v\}$ is present if there is currently a request on the cross-connect from some input fiber to some output fiber such that the request starts on wavelength u and ends on wavelength v. The color assigned to the edge represents the wavelength interchanger that the demand is routed through.

Since the inexpensive part of the cross-connect can route channels to and from the wavelength interchangers arbitrarily, the identities of the input and output fibers for a particular demand are not needed for the graph model. However, the *number* of input (or output) fibers is critical because it bounds the number of channels of a given wavelength, thus the degree of the graph.

Figure 2 shows a 2×2 cross-connect with 3 wavelength interchangers handling 4 wavelengths. Figure 3 shows the corresponding graph, with $n = 4$ nodes

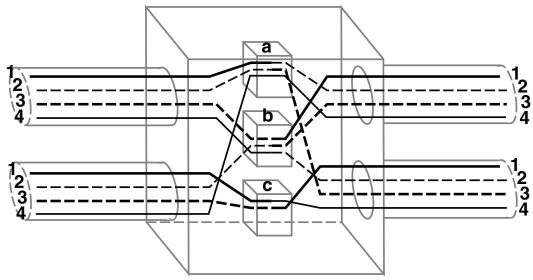

Fig. 2. A 2×2 cross-connect with 4 wavelengths and 3 wavelength interchangers

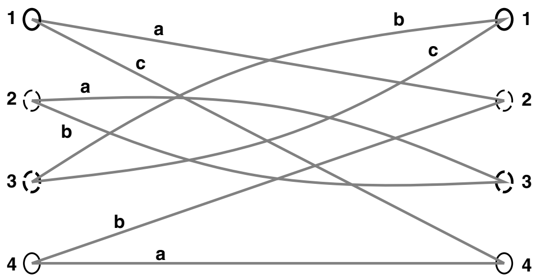

Fig. 3. The graph corresponding to the cross-connect in Figure 2

on a side and maximum degree $\Delta = 2$. We have labeled the wavelengths (and nodes) numerically, and the wavelength interchangers (and colors) by letters a, b and c. Table 1 shows a mapping of terminology and notation between WDM cross-connects and the graph model.

Obviously at least Δ colors will be necessary to edge color a dynamic graph if a node can have as many as Δ incident edges; on the other hand with $2\Delta-1$ colors the colorer can employ any strategy and will never be stymied. We show in Section 6.1 that for every edge coloring algorithm there exists a dynamic graph with only $O(\Delta^2)$ nodes that requires $2\Delta-1$ colors. This result holds also for simple graphs.

This raises the question of whether smaller dynamic graphs require fewer than $2\Delta-1$ colors. In Section 6.2 we consider bipartite dynamic graphs with very few nodes. These graphs can be edge colored with substantially fewer than $2\Delta-1$ colors. In particular, dynamic graphs with 2 nodes on each side of the bipartition can be edge colored with $3\Delta/2$ colors and this bound is tight. For bipartite dynamic graphs with 3 nodes on each side, $15\Delta/8$ colors suffice and there is a lower bound of $7\Delta/4$. Due to space constraints, we omit all proofs in Section 6.2.

This leaves open the question of how many colors are necessary when the number of nodes in the graph is between a small constant and $O(\Delta^2)$. In Section 6.3 we provide a lower bound of $2(1-\varepsilon)\Delta$ colors for graphs with at least $1/\varepsilon^2$ nodes.

Table 1. Mapping of terminology and notation of WDM cross-connects and graph model

WDM cross-connects	**graph model**
Λ = # of wavelengths	$n = \max(\lvert A\rvert, \lvert B\rvert)$
demand	edge
k = # input/output fibers	Δ = maximum node degree
# wavelength interchangers	# edge colors

4 Edge Coloring of Graphs

A proper edge coloring of a graph G requires that adjacent edges be assigned distinct colors. The minimum number of colors needed to color the edges of a graph, usually called the chromatic index, is a classical graph parameter that has been studied extensively; see e.g. [4] or [12]. König [15] showed that every bipartite graph with maximum degree Δ has chromatic index at most Δ; Vizing [22] proved that every *simple* graph with maximum degree Δ has chromatic index either Δ or $\Delta+1$. Even so, for many classes of graphs, determining the exact chromatic index is NP-complete [11, 5, 10].

More recently this problem has been considered in settings in which the entire graph is not known in advance. One such body of work considers *constrained* edge colorings in which a partially colored graph is given as input. The remaining edges must be legally colored without ever re-coloring any edges [13, 14, 7, 16, 17, 6]. In the more standard version of the on-line edge coloring problem, the graph is presented one edge (or node) at a time and each edge must be colored by the algorithm as it is presented. Favrholdt and Nielsen [8] consider on-line edge coloring with a fixed number of colors, the goal being to color as many edges as possible. Bar-Noy et al. [1] provide a graph of maximum degree Δ and $O(\binom{2\Delta-1}{\Delta})$ nodes for which any on-line edge coloring algorithm requires $2\Delta-1$ colors; this implies that for graphs with n nodes and $O(\log n)$ maximum degree one can do no better than the greedy on-line coloring algorithm in the worst-case.

5 Problem Definitions

Let $\delta_E(v)$ be the degree of node v given a set E of edges. Given a set N of nodes and a maximum degree Δ define a set of edges E to be *valid* if and only if

1. for any $\{u, v\} \in E$, $u, v \in N$;
2. for any $v \in N$, $\delta_E(v) \leq \Delta$.

Let the *edge sequence* $\mathcal{E} = (E_1, E_2, \ldots)$ be a sequence of valid edge sets E_i. We define a dynamic graph $\mathcal{G}(N, \Delta, \mathcal{E})$ to be the sequence of graphs $(G_1, G_2 \ldots)$ such that $G_t = (N, E_t)$. We assume that $\mathcal{G}$ starts as just the set N of nodes and define $E_0 = \emptyset$ to be the initial set of edges.

A coloring C assigns a color $C(e)$ to each edge e. We define $\mathcal{C} = (C_1, C_2 \ldots)$ to be a *proper coloring* of $\mathcal{G}$ if for any i

1. C_i is a proper coloring of G_i;
2. for any e such that $e \in E_{i-1}$ and $e \in E_i$, $C_{i-1}(e) = C_i(e)$.

If the node set N consists of "left nodes" from a set A and "right nodes" from a set B, and all edges connect a node from A with a node from B, then the dynamic graph is bipartite and we write $\mathcal{G}(A, B, \Delta, \mathcal{E})$ instead of $\mathcal{G}(N, \Delta, \mathcal{E})$. All of the dynamic graphs we construct here are bipartite, on account of our intended application.

Our lower bounds apply to any algorithm, deterministic or randomized. However, the dynamic graph that witnesses the lower bound may depend on past choices made by the algorithm. Thus it is possible that there exists a randomized algorithm whose *expected* performance is not governed by our lower bounds.

6 New Results

We begin by showing that for every edge coloring algorithm there exists a dynamic graph with $O(\Delta^2)$ nodes that requires $2\Delta - 1$ colors.

6.1 A Lower Bound for Polynomial Size Graphs

We show that for any edge coloring algorithm there is a dynamic bipartite graph $\mathcal{G}(A, B, \Delta, \mathcal{E})$ with $\max(|A|, |B|) = n = (\frac{1}{4} + o(1))\Delta^2$ such that the algorithm must use $2\Delta - 1$ colors to edge color $\mathcal{G}$.

Define the *spectrum* $S(v)$ of a node v to be the set of colors of its incident edges. We say that a node is *full* if its degree is Δ, i.e. $|S(v)| = \Delta$.

We begin by showing that if there is some edge sequence that at some time results in a graph with a particular property then for any edge coloring algorithm we can construct an edge sequence that will require $2\Delta - 1$ colors.

Lemma 1. *Let $\mathcal{G} = \mathcal{G}(A, B, \Delta, \mathcal{E})$ be a dynamic bipartite graph that has been colored so that at a certain stage j, $G_j(A, B, E_j)$ has $m = 1 + \lceil \log_2 \Delta \rceil$ full nodes with the same spectrum S where*

1. *one of the m points (say, x) is on one side of the bipartition, the rest $(y_1, \ldots, y_{m-1})$ are all on the other side;*
2. *there are currently no edges between x and any of the y_i.*

Then for any edge coloring algorithm there is an edge sequence $\mathcal{E}' = (E'_1, E'_2, \ldots)$ where $E'_r = E_r$ for $r = 1, 2, \ldots, j$ such that the algorithm will require 2Δ–1 colors to edge color $\mathcal{G}(A, B, \Delta, \mathcal{E}')$.

Proof. The edge sequence E' is defined for E_r, $1 \le r \le j$ and now we describe which edges are to be deleted and which edges are to be added to progress from E'_i to E'_{i+1} for $i \ge j$. Consider any edge coloring algorithm and let C'_i be the coloring of the edges of E'_i by that algorithm.

Define $S_i(v)$ to be the spectrum of node v according to C'_i. Let $T_i := S_i(x) \cap S$ and $t_i := \lfloor |T_i|/2 \rfloor$. Let X_i and Y_i be disjoint subsets of T_i of size t_i. To define

E'_{i+1} from E'_i, delete the edges of E'_i incident to x whose colors lie in X_i, and the edges of E'_i incident to y_{i-j} whose colors lie in Y_i. Notice that even after removing these edges, the union of the current spectrums of x and y_{i-j} still contains S. Replace the removed edges by t_i new edges from x to y_{i-j}. Notice that all of their colors must lie outside the set S, thus $S_{i+1}(x)$ now has t_i new colors.

Thus C'_{i+1} is such that $|T_{i+1}| = \lceil |T_i|/2 \rceil$ so that after precisely $m-1 = \lceil \log_2 \Delta \rceil$ stages we have $|T_{j+m-1}| = 1$, thus $|S_{j+m-1}(x) \cup S| = 2\Delta - 1$. The total number of nodes originally required is thus $m = 1 + \lceil \log_2 \Delta \rceil$. □

Now what remains for us to show is that for any edge coloring algorithm we can define $E_1, E_2, \ldots, E_j$ so that the edge coloring algorithm will either use all $2\Delta - 1$ colors or force the conditions stated in Lemma 1. In what follows, when we speak of edges being colored we mean by some arbitrary edge coloring algorithm.

We choose opposing nodes x and y, connect by Δ edges, and name the resulting colors 1 through Δ. These will be called the *light* colors. Colors Δ through $2\Delta - 1$ are said to be *dark*. Note that color Δ is both light and dark and thus there are at most $\Delta - 1$ colors that are only light and at most $\Delta - 1$ colors that are only dark.

Next we choose Δ new points on each side, say u_1 through u_Δ on the left and v_1 through v_Δ on the right, and connect completely to form a copy of the complete $\Delta \times \Delta$ bipartite graph, $K_{\Delta,\Delta}$. Suppose first that at least 3/4 of the colors of these Δ^2 edges are dark.

Then we proceed to construct a point with completely dark spectrum, which, in combination with the $\{x, y\}$ edges, uses all colors. To do this, we direct our attention to new nodes $w_1, w_2, \ldots, w_k$ on the right, and successively connect all the u_i's to w_1, then to w_2 etc., each time discarding the light edges and keeping the dark ones. Since there are only $\Delta - 1$ colors that are not dark and we add Δ edges to each w_i, at least one new dark edge must appear each time. Hence, after $k = \Delta^2/4$ iterations some u_i must be full and dark, and so all $2\Delta - 1$ colors are being used.

If, on the other hand, more than 1/4 of the $K_{\Delta,\Delta}$ edge colors are light, we construct instead a light point. Again we successively connect all the u_i's to $w_1, w_2, \ldots, w_k$, this time discarding the dark edges and keeping the light ones. Here, however, we gain at least 3 light edges per w_i since if w_i has only one light incident edge we already have all the colors, and if it has just two light colors then we do the following: we remove one of the light edges $\{x, y\}$ that has the same color as one of the edges incident with w_i. We delete the other edge incident with w_i having a light color. Note that the spectra of x and w_i are disjoint and together consist of $2\Delta - 2$ colors. We then add a new edge $\{x, w_i\}$ and this edge must be given a color different from the 2Δ–2 colors in the spectra of x and w_i and so all $2\Delta - 1$ colors will be used.

Hence after only $k = (\frac{1}{3})(3\Delta^2/4)$ iterations we have a light point, say u_1, among the u_i's; replace u_1 by u'_1 and continue until there are $1 + \lceil \log_2 \Delta \rceil$ light points. Each time a light point is replaced Δ light edges are lost from the current batch of u_i's and u_i''s, so altogether $\Delta^2/4 + \Delta\lceil\log_2 \Delta\rceil$ w_i's may be needed.

Either way we have $n = (\frac{1}{4} + o(1))\Delta^2$.

Combining this construction with Lemma 1 (using x or y as the "lonely" point, as necessary) allows us to conclude the following.

Theorem 1. *For any edge coloring algorithm, there is a dynamic bipartite graph* $\mathcal{G}(A, B, \Delta, \mathcal{E})$ *with* $\max(|A|, |B|) = n$ *where* $n = (\frac{1}{4} + o(1))\Delta^2$ *such that the edge coloring algorithm must use at least* 2Δ–1 *colors to color* $\mathcal{G}$*. (Thus, in particular, a* $k \times k$ *wide-sense nonblocking cross-connect in an* n*-wavelength system requires a full complement of* $2k-1$ *wavelength interchangers.)*

A similar argument gives the same result for simple graphs, however the constant factor in $n = O(\Delta^2)$ is larger than $1/4 + o(1)$.

Given the result of Theorem 1 it is natural to ask whether for all edge coloring algorithms and all n, there exists a dynamic graph of size n with maximum degree Δ that requires 2Δ–1 colors to be edge colored. The following results show that for small graphs strictly fewer than $2\Delta-1$ colors are sufficient.

6.2 Small Bipartite Graphs Need Fewer Colors

We address the case of dynamic bipartite graphs $\mathcal{G}(A, B, \Delta, \mathcal{E})$ with $n = |A| = |B|$ nodes where n is small, i.e. $n = 2$ or $n = 3$. For $n = 2$ there is an algorithm that uses at most $3\Delta/2$ colors and for $n = 3$ there is an algorithm that uses at most $15\Delta/8$ colors. We then consider lower bounds for these two cases. We assume throughout this section that Δ is divisible by 8. We omit all proofs in this section due to space constraints.

Theorem 2. *If* $\mathcal{G}(A, B, \Delta, \mathcal{E})$ *has* $|A| = |B| = 2$ *then it can be edge colored with* $3\Delta/2$ *colors.*

Theorem 3. *If* $\mathcal{G}(A, B, \Delta, \mathcal{E})$ *has* $|A| = |B| = 3$*, then it can be edge colored with* $15\Delta/8$ *colors.*

The next results show that our upper bound on the number of colors sufficient to edge color 2×2 dynamic bipartite graphs is tight whereas there remains a gap for 3×3 dynamic bipartite graphs.

Theorem 4. *For any edge coloring algorithm, there exists a dynamic bipartite graph* $\mathcal{G}(A, B, \Delta, \mathcal{E})$ *where* $|A| = |B| = 2$ *for which the edge coloring requires at least* $3\Delta/2$ *colors.*

Theorem 5. *For any edge coloring algorithm, there exists a dynamic bipartite graph* $\mathcal{G}(A, B, \Delta, \mathcal{E})$ *where* $|A| = |B| = 3$ *such that the edge coloring algorithm must use at least* $7\Delta/4$ *colors.*

The results of this section lead to the following question: How many colors are necessary and sufficient to edge color all dynamic graphs with maximum degree Δ and $o(\Delta^2)$ nodes?

6.3 A Lower Bound

Theorem 6. *For any edge coloring algorithm, any $\varepsilon > 0$ and $\Delta > 1/2\varepsilon$, there exists a dynamic bipartite graph with fewer than $1/\varepsilon^2$ nodes on each side that requires the algorithm to use more than $2(1-\varepsilon)\Delta$ colors.*

Note that the lower bound on Δ is necessary, since if $\Delta \leq 1/2\varepsilon$ then $2(1-\varepsilon)\Delta \geq 2\Delta-1$, and we can never force more than $2\Delta-1$ colors.

Proof. In view of Theorems 4 and 5 we may assume $\varepsilon < \frac{1}{8}$, thus $\Delta > 4$ and $2(1-\varepsilon)\Delta > 2(1-\frac{1}{8})4 > 7$. Let us fix ε and Δ accordingly, put $q = \lceil \varepsilon\Delta \rceil$ and let C be a set of $\lfloor 2(1-\varepsilon)\Delta \rfloor$ colors. We will construct a dynamic graph whose edge coloring from C leads to a contradiction. The method employs a variation of part of the proof of Theorem 1, and indeed by letting Δ be a function of ε one could deduce a weaker form of that theorem.

We consider first the case in which $\lceil 2\varepsilon\Delta \rceil$ is even, thus equal to $2q$; then $|C| = 2\Delta-2q$. Let $m = \lceil 1/\varepsilon \rceil < 8\varepsilon/7$. Let $X := \{x_1, \ldots, x_m\}$ be a set of m left nodes and $Y := \{y_1, \ldots, y_m\}$ a set of right nodes, each x_i matched by Δ parallel edges to y_i. When these edges have been colored we choose $A \subset C$ with $|A| = |C|/2 = \Delta-q$ so as to maximize the number of edges colored by A. All edges *not* colored from A are now removed.

Let s be the sum of the degrees of the nodes in X, so that $s \geq m\Delta|A|/|C| \geq m\Delta/2$ by choice of A. Since no node can have more than $\Delta-q$ incident edges colored by A, we also have $s \leq m(\Delta-q) \leq m\Delta - \Delta$; thus there are at least Δ places for new edges to be introduced, incident to nodes in X.

Next we consider a new node z_1 on the right, adding a full complement of Δ edges between z_1 and X and then deleting all edges that did not get colored by colors in A. Since at least $\Delta - (|C|-|A|) = q$ of the edges incident to z_1 must have been A-colored, the degree sum s will increase by at least q.

We now repeat the operation with more right-hand nodes $z_2, z_3, \ldots, z_t$ where $t = \lceil 4/7\varepsilon^2 \rceil \geq m/2\varepsilon$. Then

$$\begin{aligned} s &\geq m\Delta/2 + tq \\ &\geq m\Delta/2 + \frac{m}{2\varepsilon}\varepsilon\Delta \\ &\geq m\Delta > m\Delta - \Delta \ , \end{aligned}$$

an impossibility.

When $\lceil 2\varepsilon\Delta \rceil$ is odd, thus equal to $2q-1$, we have to give away a bit more. Then $|C| = 2\Delta - 2q + 1$; and $q \geq 2$ since by assumption $2\varepsilon\Delta > 1$, and it follows that $q-1 > \frac{2}{3}\varepsilon\Delta$.

Let $m = \lceil 3/2\varepsilon \rceil \leq \frac{13}{8}\varepsilon$, and select A as above but with $|A| = \Delta-q+1 > |C|/2$. We now have $s > m\Delta/2$ and $s \leq m|A| = m\Delta - m(q-1) \leq m\Delta - (3/2\varepsilon)(\frac{2}{3}\varepsilon\Delta) \leq m\Delta - \Delta$, so again there are at least Δ places for new edges to be introduced.

As before $\Delta - (|C|-|A|) = q$ of the edges incident to each successive z_i must be colored by A, so the degree sum increases by at least q each time. Taking $t = 13/16\varepsilon^2 > 2m/\varepsilon$ now causes the contradiction in the same manner as above. $\square$

A simpler version of the proof of Theorem 6 above shows that when Δ is large relative to $1/\varepsilon$, we get a stronger but asymptotic result, namely that $(1 + o(1))/2\varepsilon^2$ nodes on a side suffice to force more than $2(1-\varepsilon)\Delta$ colors.

A similar construction shows that for any edge coloring algorithm, there is a simple dynamic bipartite graph with $|A| = |B| \geq \max\left[2(1-\varepsilon)\Delta, (1 + \frac{(1-\varepsilon)}{\varepsilon}\Delta)\right]$ that requires the algorithm to use $2(1-\varepsilon)\Delta$ colors.

7 Conclusions and Future Work

We have presented a variety of results concerning the number of wavelength interchangers needed in a wide-sense nonblocking $k \times k$ WDM cross-connect by considering the problem of edge coloring dynamic graphs with maximum degree $\Delta = k$. In particular, for the case of 2 wavelengths, the necessary and sufficient number of wavelength interchangers is $3k/2$. When there are 3 wavelengths, a lower bound of $7k/4$ and an upper bound of $15k/8$ was given. However we showed that if there are about $k^2/4$ or more wavelengths then $2k-1$ wavelength interchangers are necessary. Thus in this case, the greedy algorithm is optimal. Furthermore, this implies that weakening the nonblocking capability from strictly nonblocking to wide-sense nonblocking does not reduce the number of wavelength interchangers needed and hence does not reduce the cost of the cross-connect. We have also shown that for any $\varepsilon > 0$ and $k > 1/2\varepsilon$, if there are at least $1/\varepsilon^2$ wavelengths then $2(1-\varepsilon)k$ wavelength interchangers are necessary.

The major remaining question would be to determine the number of wavelength interchangers necessary and sufficient for such cross-connects supporting $o(k^2)$ wavelengths. In particular, it would be interesting to know the smallest number of wavelengths so that any wide-sense nonblocking $k \times k$ WDM cross-connect supporting these wavelengths would require $2k-1$ wavelength interchangers. We would also like to know whether it is true that for every fixed number of wavelengths, say Λ, there is some $c < 2$ such that there is a wide-sense nonblocking $k \times k$ WDM cross-connect supporting Λ wavelengths with only ck wavelength interchangers.

Along the lines of [18] one could consider using less powerful wavelength interchangers (e.g. those with the ability to swap only two wavelength channels leaving the others fixed) and ask how using such weaker wavelength interchangers affects the number of wavelength interchangers needed in a wide-sense nonblocking cross-connect.

References

[1] A. Bar-Noy, R. Motwani, and J. Naor. The greedy algorithm is optimal for on-line edge coloring. *Information Processing Letters*, 44(5):251–253, 1992. 543

[2] L. A. Bassalygo and M. S. Pinsker. Complexity of an optimal nonblocking switching network without reconnections. *Problems Inform. Transmission*, 9:64–66, 1974. 541

[3] 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. 541

[4] C. Berge. *Graphs and Hypergraphs*. North Holland, Amsterdam, 1973. 543

[5] L. Cai and J. A. Ellis. NP-completeness of edge-coloring some restricted graphs. *Discrete Appl. Math.*, 30:15–27, 1991. 543

[6] I. Caragiannis, C. Kaklamanis, and P. Persiano. Edge coloring of bipartite graphs with constraints. *Theoretical Computer Science*, 270(1-2):361–399, 2002. 543

[7] T. Erlebach, K. Jansen, C. Kaklamanis, M. Mihail, and P. Persiano. Optimal wavelength routing on directed fiber trees. *Theoretical Computer Science*, 221(1-2):119–137, 1999. 539, 543

[8] L. Favrholdt and M. Nielsen. On-line edge coloring with a fixed number of colors. In *Foundations of Software Technology and Theoretical Computer Science*, pages 106–116, Dec. 2000. 543

[9] P. Feldman, J. Friedman, and N. Pippenger. Wide-sense nonblocking networks. *SIAM J. Disc. Math.*, 1(2):158–173, 1988. 541

[10] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Co., San Francisco, CA, 1979. 543

[11] I. Holyer. The NP-completeness of edge-coloring. *SIAM Journal of Computing*, 10:718–720, 1981. 543

[12] T. R. Jensen and B. Toft. *Graph Colouring Problems*. Wiley, New York, 1995. 543

[13] C. Kaklamanis and P. Persiano. Efficient wavelength routing on directed fiber trees. In *Proceedings of the 4th European Symposium on Algorithms (ESA '96)*, pages 460–470, 1996. 539, 543

[14] C. Kaklamanis, P. Persiano, T. Erlebach, and K. Jansen. Constrained bipartite edge coloring with applications to wavelength routing. In *Proceedings of the 24th International Colloquium on Automata, Languages, and Programming (ICALP '97)*, pages 493–504, 1997. 539, 543

[15] D. König. Graphok és alkalmazásuk a determinánsok és a halmazok elméletére (in Hungarian). *Mathematikai és Természettudományi Értesítő*, 34:104–119, 1916. 543

[16] V. Kumar and E. J. Schwabe. Improved access to optical bandwidth. In *Proceedings of 8th ACM-SIAM Symposium on Discrete Algorithms(SODA '97)*, pages 437–444, 1997. 539, 543

[17] M. Mihail, C. Kaklamanis, and S. Rao. Efficient access to optical bandwidth. In *Proceedings of the 36th Annual IEEE Symposium on the Foundations of Computer Science (FOCS '95)*, pages 548–557, 1995. 539, 543

[18] R. Ramaswami and G. H. Sasaki. Multiwavelength optical networks with limited wavelength conversion. In *Proceedings of IEEE INFOCOM*, volume 2, pages 489–498, 1997. 548

[19] A. Rasala and G. Wilfong. Strictly non-blocking WDM cross-connects. In *Proceedings of Symposium on Discrete Algorithms (SODA '00)*, pages 606–615, 2000. 540, 541

[20] A. Rasala and G. Wilfong. Strictly non-blocking WDM cross-connects for heterogreous networks. In *Proceedings of Symoposium on Theory of Computation (STOC '00)*, 2000. 540

[21] C. E. Shannon. Memory requirements in a telephone exchange. *Bell System Tech. J.*, 29:343–349, 1950. 540

[22] V. G. Vizing. On an estimate of the chromatic class of a p-graph (in Russian). *Diskret. Analiz*, 3:23–30, 1964. 543

[23] 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. 539, 541