Bandwidth Allocation Algorithms on Tree–Shaped All–Optical Networks with Wavelength Converters*

Vincenzo Auletta,[1] Ioannis Caragiannis,[2] Christos Kaklamanis,[2,3] and Pino Persiano[1]

[1] Dipartimento di Informatica ed Applicazioni, Universita di Salerno, 84081 Baronissi (SA), Italy. (e-mail {auletta,persiano}@dia.unisa.it)
[2] Computer Technology Institute, Dept. of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece. (e-mail: {kakl,caragian}@ceid.upatras.gr)
[3] Computer Technology Institute, 26500 Rio, Greece. (e-mail:kakl@cti.gr)

Abstract. In this paper we consider the problem of allocating optical bandwidth on a tree-shaped network when wavelength conversion is allowed. We consider both converters with full conversion capabilities and converters with limited capabilities.

We give upper and lower bounds on the number of full wavelength converters necessary to fully utilize the optical bandwidth available.

For converters of limited capabilities we show that, for general trees, converters of degree $\lceil L/3 \rceil + 1$ are sufficient to route a set of requests of maximum load L using at most $3/2L$ wavelengths. For the special case of binary trees, we show that converters of degree $O(\sqrt{L})$ are sufficient to route a set of requests of maximum load L using exactly L wavelengths. The study of limited capabilities converters is based on graph construction problems that might be of their own interest.

1 Introduction

1.1 Background

Optical fiber is rapidly becoming the standard transmission medium for networks, and can provide the required data rate, error rate and delay performance for future networks. However, data rates are limited in opto–electronic networks by the need to convert the optical signals on the fiber to electronic signals in order to process them at the network nodes. Electronic parallel processing techniques are capable, in principle, to meet future high data rate requirements, but the opto–electronic conversion is expensive. It appears likely that, as optical technology improves, simple optical processing will remove the need for electro–optic conversion. Networks using optical transmission and maintaining optical data paths through the nodes are called all–optical networks.

* Research supported by Progetto MURST 40%, Algoritmi, Modelli di Calcolo e Strutture Informative and by ESPRIT Project 20244 ALCOM–IT.

Optical technology is not as mature as conventional technology. There are limits in how sophisticated optical processing at each node can be done. In an attempt to keep the processing inside the all–optical network simple, most of the work in all–optical networks has been limited to broadcast and select networks [10].

Our network architectural approach is based on the use of Wavelength Division Multiplexing (WDM), the exploitation of available wavelengths to route the signal to its intended destination (wavelength routing), the use of wavelength–selective switches, and the translation of signals from one wavelength to another, and ultimately, does not require synchronization and central control. These principles permit networks to be built whose size is essentially unlimited and are therefore suitable for Wide Area Networks (WAN's). The basic building blocks of our approach form integrated WDM transmitter and receiver arrays, wavelength–selective switches, wavelength converters and collision detected couplers.

The WDM technology establishes connectivity by finding transmitter–receiver paths and assigning a wavelength to each path such that no two paths going through the same link use the same wavelength. Optical bandwidth is the number of available wavelengths. The optical bandwidth is a scarce resource. State–of–the–art technology allows no more than 100 wavelengths in the laboratory, less than half in manufacturing, and there is no anticipation of dramatic progress in the near future.

Current techniques for optical bandwidth allocation cannot guarantee a high degree of bandwidth utilization under the worst conditions. In particular, this factor is 50% for ring networks, 60% for trees. A promising solution for efficient use of bandwidth is wavelength conversion. Devices called wavelength converters are located at the nodes of the network and they can change the wavelength assigned to a transmitter–receiver path up to a node and allocate a different wavelength at the rest of the path. In this way, an efficient bandwidth allocation can be achieved that resolves wavelength conflicts and improves bandwidth utilization.

In each node, wavelength conversion can be done either all–optically, or using opto–electronic conversion. Since our purpose is to improve bandwidth allocation in all–optical networks, we concentrate on the use of all–optical wavelength converters. Such devices are expected to have limited conversion capabilities. The disadvantage of their use is increased cost and complexity.

1.2 Related Work

Several authors have already address the case of no wavelength conversion. We point out the work of Pankaj [15] who considered uniform loads on shuffle exchange, De Bruijn, and hypercube networks, and the papers of Aggarwal et al. [1] and Raghavan and Upfal [16] who, among other issues, obtained bounds for uniform loads on meshes, bounds in terms of expansion, and arbitrary load bounds for specific topologies such as trees, rings, and trees of rings. Raghavan and Upfal [16] showed that routing requests of maximum load L per link of undirected trees can be satisfied using $3L/2$ optical wavelengths and their arguments extend to give a $2L$ bound for the directed case.

Mihail *et al.* [13] addressed the directed case. Their main result is a $15L/8$ upper bound. This is done by reduction to a bipartite graph edge-coloring, which is achieved in phases by obtaining matchings of the bipartite graph, and coloring them in pairs using detailed potential and averaging arguments. Kaklamanis and Persiano [8] (and independently Kumar and Schwabe [11]) improved the upper bound for directed trees to $7L/4$. The main idea of their algorithms is similar to the one of [13] but new techniques are used for partitioning the bipartite graph matchings into groups that can be colored and accounted for independently. Kaklamanis *et al.* [9] have obtained a greedy algorithm that routes a set of requests of maximum load L using at most $5L/3$ wavelengths and have proved that no greedy algorithm can go below $5L/3$ in general. Caragiannis *et al.* [3] and Jansen [7] present simple greedy algorithms that meets the optimal ratio of $5/3$ for binary trees. In [3], it is also proved that the lower bound of $5L/3$ holds even for the case of leaf–to–leaf communication.

Variants of the limited conversion model are considered by Ramaswami and Sasaki [17], Subramaniam et al. [18], Yates et al. [19], and Lee and Li [12]. The work in [18] studies sparse wavelength conversion, where networks comprise a mix of nodes having either full or no wavelength conversion. The wavelength assignments in [18, 19] are simple heuristics, and their performance analyses are based upon probabilistic models and techniques. Ramaswami and Sasaki [17] propose ring and star networks with limited wavelength conversion to support sets of lightpath efficiently. In particular, they consider L–wavelength technology and give a ring network with one node having fixed wavelength conversion and the rest of the nodes with no wavelength conversion such that all patterns of requests with maximum load $\lambda_{max} \leq L-1$ have a legal wavelength assignment. They also present a ring network with two nodes with wavelength degree two and the rest of the nodes with no wavelength conversion such that all patterns of requests with maximum load $\lambda_{max} \leq L$ have a legal wavelength assignment. Although they address the undirected case, all their results for rings can be applied to the directed case. Furthermore, they propose algorithms for bandwidth allocation in undirected stars, trees, and networks with arbitrary topologies where route lengths are at most two. An all–optical wavelength converter based on four–wave mixing is presented in [20].

1.3 Network Model

We model the underlying fiber network as a directed graph. Connectivity requests are ordered pairs of nodes, to be thought of as transmitter–receiver pairs. For networks with unique transmitter–receiver paths (such as trees), the load of a fiber link is the number of paths going through the link. WDM technology establishes connectivity by finding transmitter–receiver paths, and assigning a wavelength to each path, so that no two paths going through the same link use the same wavelength. Intuitively we can think of the wavelengths as colors. Optical bandwidth is the number of available wavelengths. Each directed fiber link can support L wavelengths $\{\omega_0, \omega_1, \ldots, \omega_{L-1}\}$, with distinct optical frequencies.

Current approaches of the wavelength assigning problem in trees use greedy algorithms [3, 8, 9, 11, 13]. A greedy algorithm visits the network in a top to bottom manner and at each vertex v colors all requests that touch vertex v and are still uncolored. Moreover, once a request has been colored it is never recolored again. Greedy algorithms are important as they are very simple and, more importantly, they are amenable of being implemented in a distributed environment.

It has been proved that greedy algorithms for bandwidth allocation in all–optical networks cannot go below the ratio of $5/3$ [9] resulting to 60% utilization of available bandwidth. Furthermore, even if a better algorithm is discovered, there exists a pattern of communication requests that requires $5L/4$ wavelengths [11] meaning that 20% of the available bandwidth across the fiber links will be unutilized. Furthermore, a modified greedy algorithm in which each node has knowledge of k-levels (instead of local knowledge) for paths that touch the node, where k is fixed, does not work since the proof of the lower bound in [9] implies that the adversary can choose the 'bad' pattern to split after $k+1$ levels. Thus a 'global' algorithm is necessary. Unfortunately, the lower bound of $5L/4$ for any algorithm shows that 100% bandwidth utilization is infeasible even for binary trees.

In order to efficiently access optical bandwidth using pure WDM technology, one can follow to approaches.

- Use a non–greedy (global) algorithm. In this case the best we can hope for is an 80% utilization of the available bandwidth [11].
- Consider the use of wavelength converters. These are devices that are able to convert wavelengths so that the same communication request can be assigned different wavelengths in different segments of the path from source to destination.

Wavelength conversion requires devices called wavelength converters located in some nodes of the network. A wavelength converter can be modelled as bipartite graph $G = (V, U, E)$. The set of vertices V and U have each one vertex for each wavelength and there is an edge between vertices $v \in V$ and $u \in U$ if and only if the converter is capable of converting the wavelength corresponding to v to the wavelength corresponding to u. For example, a full converter corresponds to a complete bipartite graph and a fixed conversion converter corresponds to a a bipartite graph where the vertices of V have degree one. Some examples of wavelength converters are depicted in Figure 1.

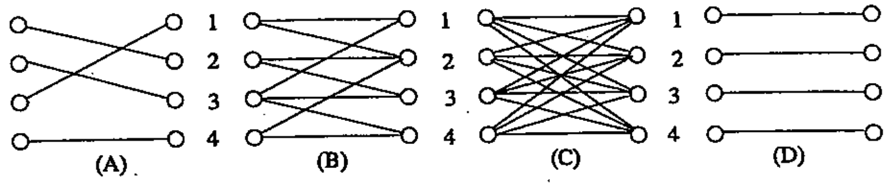

Fig. 1. Wavelength Conversion Bipartite Graphs: (A) Fixed conversion, (B) Partial Conversion, (C) Full conversion, and (D) No conversion.

Definition 1. The degree of a converter is the maximum degree of the vertices $v \in V$ of its corresponding bipartite graph.

We can define the class of greedy algorithms for networks that support wavelength conversion. Such algorithms visit the network in a DFS manner but the functionality in a node v that supports wavelength conversion is different. In this case, the greedy algorithm can assign a different color for a request that has been colored in a previous step, so that the wavelength corresponding to the old color can be converted to the new one by the wavelength converter located in node v. For a request that touches node v and has not been colored, the algorithm can assign a pair of wavelengths (colors) ω_i, ω_j since wavelength ω_i can be converted to ω_j (or ω_j can be converted to ω_i) by the converter of node v. Throughout this work, converters of identical functionality are located at nodes of the network.

1.4 Outline of the Paper

In this paper, we consider both tree networks with a small number of complex full conversion wavelength converters or networks where on each link there is a wavelength converters of small degree. We concentrate on the use of greedy algorithms and derive several upper and lower bounds. Our results can be summarized as follows.

- A lower bound on the number of full wavelength converters necessary to route a pattern of communication requests with maximum load L per directed fiber link on a tree network using at most L wavelengths.

- An upper bound on the number of full wavelength converters sufficient to route a pattern of communication requests with maximum load L per directed fiber link on a tree network using at most L wavelengths.

- A greedy algorithm that routes a set of requests of maximum load L using $3L/2$ wavelengths and converters of degree $\lceil L/3 \rceil + 1$. This beats the $5/3L$ lower bound for greedy algorithm of [9].

- A greedy algorithm that routes a set of requests of maximum load L using L wavelengths and converters of degree $O(L^{1/2})$ for binary trees. This beats the $5/4L$ lower bound on the number of wavelength necessary for routing sets of requests of maximum load L on binary trees [11].

2 A Lower Bound on the Number of Wavelength Converters

In this section we study the number of full converters necessary in a tree to route any set of of requests of load at most L using less than $5/4L$ wavelengths.

Theorem 1. *For each L, there exists a tree network with N nodes which has wavelength converters at $\lfloor \frac{N}{4} - \frac{3}{2} \rfloor$ nodes, and a pattern of communication requests R with maximum load L per directed fiber link, such that no algorithm can assign less than $5L/4$ wavelengths.*

Proof. Consider the tree network $T_n(V_n, E_n)$ where $V_n = \{v_i | i = 0, 1, ..., n\} \cup \{u_i | i = 1, 2, ..., n-1\}$ and $E_n = \{(v_i, v_{i+1}) | i = 0, 1, ..., n-1\} \cup \{(v_i, u_i) | i = 1, 2, ..., n-1\}, n \geq 3$. For this network, we define the pattern of communication requests:

$$\begin{aligned} R_n = & \{(u_i, u_{i+2}) | i = 1, 2, ..., n-2\} \cup \{(u_{i+1}, u_i) | i = 1, 2, ..., n-2\} \\ \cup & \{(v_0, u_1), (v_0, u_2), (u_1, v_0), (u_{n-2}, v_n), (u_{n-1}, v_n), (v_n, u_{n-1}), (v_n, v_0)\} \end{aligned}$$

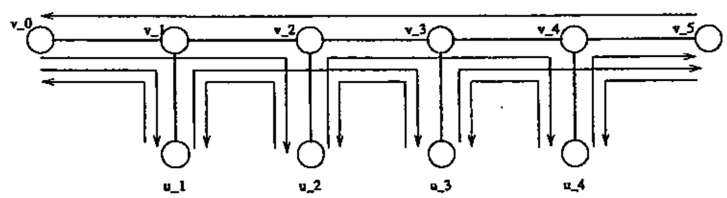

Fig. 2. The network T_n with the pattern of communication requests R_n.

This network is depicted in Figure 2. Such networks have some interesting properties. For $0 \leq j, k \leq n$ and $k - j \geq 3$, any subgraph $H_{j,k}(V_{j,k}, E_{j,k})$ of T_n, such that $V_{j,k} = \{v_i | i = j, ..., k\} \cup \{u_i | i = j+1, ..., k-1\}, E_{j,k} = \{(v_i, v_{i+1}) | i = j, ..., k-1\} \cup \{(v_i, u_i) | i = j+1, ..., k-1\}$ and the corresponding pattern of communication requests:

$$\begin{aligned} Q_{j,k} = & \{(u_i, u_{i+2}) | i = j+1, ..., k-3\} \cup \{(u_{i+1}, u_i) | i = j+1, 2, ..., k-2\} \\ \cup & \{(v_j, u_{j+1}), (v_j, u_{j+2}), (u_{j+1}, v_j), (u_{k-2}, v_k), (u_{k-1}, v_k), (v_k, u_{k-1}), (v_k, v_j)\} \end{aligned}$$

is a tree network T_{k-j} with a pattern of communication requests R_{k-j}.

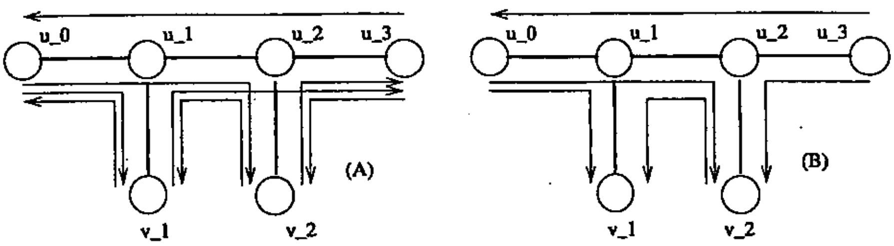

Fig. 3. Simple networks with patterns of communication requests that require $5L/4$ wavelengths.

Furthermore, the network T_3 with the pattern of communication requests R_3 (Figure 3A) is similar to that used by Kumar and Schwabe [11] to show a lower bound of $5L/4$ (Figure 3B). The pattern R_3 is in fact a superset of the pattern mentioned in their paper. We will use the generalized result as a lemma.

Lemma 1. *For any $n \geq 3$ and $L \geq 2$, the pattern R_n with (maximum) load L per directed fiber link in the tree network T_n, requires at least $5L/4$ wavelengths.*

In order to efficiently allocate the optical bandwidth in such a network, we must locate C wavelength converters in some nodes $v_{k_1}, v_{k_2}, \ldots, v_{k_C}$ such that:

- the subgraph $H_{k_i,k_{i+1}}$, is not a $T_{k_{i+1}-k_i}$ network with a pattern of communication requests $R_{k_{i+1}-k_i}$, for $1 \leq i \leq C-1$,
- the subgraph H_{0,k_1} is not a T_{k_1} network with a pattern of communication requests R_{k_1} and
- the subgraph $H_{k_C,n}$ is not a T_{n-k_C} network with a pattern of communication requests R_{n-k_C}.

This cannot be achieved unless we locate wavelength converters in such a way that any node without wavelength conversion is adjacent to two nodes with wavelength conversion. Thus, for n even, at least $\frac{n}{2}-1$ wavelength converters must be located in nodes $v_2, v_4, \ldots, v_{n-2}$ and for odd values of n, at least $\frac{n-1}{2}$ wavelength converters must be located either in nodes $v_1, v_3, \ldots, v_{n-2}$ or in nodes $v_2, v_4, \ldots, v_{n-1}$. The total number of nodes in a T_n network is $N = |V| = 2n$ and the minimum number of wavelength converters is exactly $C = \lfloor \frac{N}{4} - \frac{1}{2} \rfloor$.

Otherwise, using (strictly) less than $\lfloor \frac{N}{4} - \frac{1}{2} \rfloor$ wavelength converters, there exists a part of the tree between two wavelength converters that constitutes a network T_m with a pattern of communication requests R_m with $m \geq 3$. In this case, $5L/4$ wavelengths are necessary.

3 An Upper Bound on the Number of Wavelength Converters

Our algorithm locates the wavelength converters as follows. Let v_0 be any node of the tree network $T(V,E)$, $|V| = N$, with degree at least 3. We perform a BFS search in the tree network starting from v_0. Each node except of the leaves and the nodes of degree 2 are assigned a integer label which indicates the distance of the node from v_0 without counting nodes with degree 2. The algorithm will locate wavelength converters either to nodes with even labels, or to nodes with odd labels. In this way, we obtain a scheme in which a node without wavelength converter is adjacent to nodes with wavelength converters and (possibly) to nodes with degree 2.

Theorem 2. *For each L, and for each tree network with N nodes it is possible to place $\frac{N}{4} - \frac{1}{2}$ full converters at the vertices of the tree so that each set of communication requests R with maximum load L per directed fiber link, can be routed using at most L wavelengths.*

Proof. Let π be the number of nodes that take part in the labeling process, Δ the number of nodes with degree 2 and Λ the number of the leaves. It is $\Lambda + 2\Delta + 3\pi \leq |E| = 2N - 2$ and $\Lambda + \Delta + \pi = |V| = N$, thus $\pi \leq \frac{N}{2} - 1 - \frac{\Delta}{2}$. The algorithm will assign converters either to the subset of nodes with odd labels or to the subset of nodes with even labels, choosing the smallest one. So, the number of converters will be $C \leq \frac{N}{4} - \frac{1}{2} - \frac{\Delta}{4}$.

Since wavelength converters split the network to stars, the idea of the algorithm is to formulate the bipartite graph [13, 8, 11, 9] at each step and color its edges with exactly L colors by finding L perfect matchings (we can do so since it the bipartite graph is L–regular). This can be performed in nodes without wavelength converters. Nodes with wavelength converters simply set the appropriate pairs of wavelengths that must be converted.

The scheme outlined in the proof of the theorem above can be extended to work also with converters of wavelength degree $L/2$. Details are omitted.

4 Limited Wavelength Conversion in Arbitrary Trees

In this section we consider wavelength converters with limited capabilities. Optical networks with limited wavelength conversion will be less costly to implement than networks wits full wavelength conversion (i.e., each wavelength can be converted to any other wavelength) but, as we shall, they provide enough conversions to use the bandwidth efficiently.

Kaklamanis *et al.* [9] give a greedy algorithm to allocate wavelengths to a set of requests of load L using at most $5/3L$ wavelengths. A greedy algorithm visits the vertices of the tree following a DFS visit. While at vertex v, the algorithm assumes that all requests touching vertices already visited have been assigned a wavelength (these are requests that go through the directed links between v and its parent) and assigns a wavelength (or a color) to all the requests touching v that have not been colored yet. In [9], it is shown that if the $2L$ requests on the directed links between v and its parent are colored with at most αL different colors, then it is possible to color the remaining requests touching v using at most $(1+\alpha/2)L$ colors in such a way that at most $\max\{\alpha L, (1+\alpha/4)L\}$ colors appear on a pair of directed links between v and any of its children. The $5/3L$ algorithm is obtained by considering the case $\alpha = 4/3L$.

As a consequence, we obtain that if we guarantee that the requests between a vertex v and its parent are colored using only L colors (i.e., $\alpha = 1$) then the remaining requests can be colored using at most $3/2L$ colors. The coloring of [9] however guarantees only that at most $5/4L$ colors are seen between v and each of its children and not L colors as we assumed for v. To reestablish the inductive hypothesis, we employ a wavelength converter for each pair of direct links. The converter between vertices u and v has to recolor the $2L$ requests going through the two directed links (u,v) using only colors $1 \ldots L$.

This leads to the following graph construction problem.

Graph Construction Problem: For each $L > 0$ and for a given constant $n \geq 1$, construct a bipartite graph $G_L = (V_1, V_2, E)$ with $|V_1| = (1+\frac{1}{n})L$, $|V_2| = L$ such that for each subset S of L vertices of V_1 there exists a perfect matching between S and V_2. We call such a family of graphs a n-rearranging family of bipartite graphs.

We next show how to construct a 2-rearranging family by giving, for each $L > 0$, a graph G_L of maximum degree $\lceil L/3 \rceil + 1$. This implies that with converters of degree $\lceil L/3 \rceil + 1$ it is possible to route sets of request of load L using $3/2L$ wavelengths. We also prove that an n-rearranging family has maximum degree no smaller than $\lceil L/(n+1) \rceil$. Thus, our construction of 3-rearranging family is indeed optimal with respect to the maximum degree. Furthermore it is possible to show that for each n there exists a n-rearranging family of maximum degree no greater than $\lceil L/(n+1) \rceil$ (the general construction is omitted from this abstract).

4.1 A 2-rearranging family of graphs

For each $L > 0$ and consider the graph $G_L = (V_1, V_2, E)$ with $|V_1| = 3/2L$ and $|V_2| = L$ and edges defined in the following way. Partition the vertices of V_1 into 3 groups denoted A_0, A_1, A_2 of size $\lceil L/2 \rceil$ each and the vertices of V_2 into 3 groups B_0, B_1, B_2 of size $\lceil L/3 \rceil$. For each $0 \leq i \leq 2$, connect each vertex of A_i to each vertex of B_i. Furthermore, for each i, each of vertex of B_i is connected to a different vertex of $A_{i+1 \bmod 3}$.

Lemma 2. *The family $\{G_L\}_{L>0}$ of graphs is a 2-rearranging family of graphs.*

Proof. We show that each subset S of L vertices of V_1 has a perfect matching with V_2 by giving an algorithm for constructing such a perfect matching.

Let S be fixed and set $S_i = S \cap A_i$ and $p_i = |S_i| - L/3$, for $i = 0, 1, 2$. Without loss of generality, suppose that p_0 and p_1 are positive and thus $p_4 = -(p_0 + p_1)$ (the other cases are similar).

Vertices of S_0 are matched to vertices of B_0 and B_1 in the following way. Choose among the vertices of S_0, p_0 vertices that have a neighbour in B_1. This is always possible as $L/3$ vertices of A_0 are connected to $L/3$ different vertices of B_1; among these, at most $|A_0| - |S_0| = L/6 - p_0$ do not appear in S_0 which implies that at least $L/2 - L/6 + p_0 = L/3 + p_0 > p_0$ vertices of S_0 have a neighbour in B_1. Each of the chosen p_0 vertices is matched to its neighbour in B_1 and the remaining vertices of S_0 to the vertices of B_0.

Now consider the vertices of S_1: all but $p_0 + p_1$ vertices can be matched with vertices in B_1 (remember p_0 vertices of B_1 have been taken by vertices of S_0). As before, we can show that there are at least $L/3 + p_1$ vertices in S_1 that have a neighbour in B_2 and, since $L/3 + p_1 \geq p_0 + p_1$, we can conclude that it is possible to match the $p_0 + p_1$ vertices of S_1 in excess to their neighbours in B_2 and the remaining vertices of S_1 to the vertices of B_1 that were not matched to vertices of S_0.

The $L/3 + p_2$ vertices of S_2 are finally matched to the $L/3 - (p_0 + p_1)$ vertices of B_2 that have not been taken yet (remember $p_2 = -(p_0 + p_1)$).

We thus have the following theorem.

Theorem 3. *Let T be a tree network with N nodes where each pair of directed links has a wavelength converter of degree $\lceil L/3 \rceil + 1$. Then any set of request of maximum load L can be routed greedily using at most $3/2L$ wavelengths.*

The above construction can be generalized to arbitrary n.

Theorem 4. *For each positive integer n, there exists a n-rearranging family of graphs $\{G_L^n\}_{L>0}$ where G_L^n has maximum degree $\lceil \frac{1}{n+1} L \rceil + 1$.*

Proof. Omitted.

Next we prove a lower bound on the degree of a n-rearranging family of graphs.

Theorem 5. *For each $L > 0$ the graph G_L belonging to a n-rearranging family has maximum degree greater than $\lceil \frac{1}{n+1} L \rceil$.*

Proof. We prove the theorem by a counting argument. Set $\alpha = 1 + 1/n$ and let x be a vertex of V_2. By hypothesis for each subset S of L vertices of V_1 there is a vertex of S that is adjacent to x. Thus, it is easily seen that x has at least $(\alpha - 1)L$ edges incident on it (details are omitted). Since we have made no particular hypothesis on vertex x we have that the same bound holds for the number of edges incident on each vertex of V_2 and the total number of edges of G_L is at least $(\alpha - 1)L^2$.

On the other hand, each edge of G_L is incident to a vertex of V_1. Thus, at least one of the αL vertices of V_1 has $\lceil \frac{\alpha - 1}{\alpha} L \rceil = \lceil \frac{1}{n+1} L \rceil$ or more edges incident on it.

5 Optimal utilization of bandwidth on binary trees

In this section we restrict our attention to greedy algorithms for allocating the optical bandwidth on directed binary tree networks. In particular, we give a simple greedy algorithm that allocates any pattern of communication requests of maximum load L per directed fiber link using only L wavelengths and converters of degree $O(\sqrt{L})$.

Consider the set of requests of Figure 4. If the requests $v \rightarrow p$ and $p \rightarrow w$ are assigned the same wavelength then the request $v \rightarrow w$ can be assigned a second wavelength, thus resulting in an optimal assignment (i.e., the number of wavelengths is equal to the load). On the other hand, if $v \rightarrow p$ and $p \rightarrow w$ are assigned different wavelengths then a third wavelength is needed for the request $v \rightarrow w$. Unfortunately, this situation can occur if a greedy algorithm (as those in [8, 9]) is used for the assignment. Thus, the idea is to convert the wavelengths assigned to the requests $v \rightarrow p$ and $p \rightarrow w$ when they traverse vertex u so that they use the same set of wavelengths when traverse directed links (v,u) and (u,w), respectively. This will allow the coloring of the requests $u \rightarrow w$ using in total at most L colors. Similarly, we would like to convert in u the wavelengths assigned to the requests $p \rightarrow v$ and $w \rightarrow p$ to color the requests $w \rightarrow v$ using at most L colors.

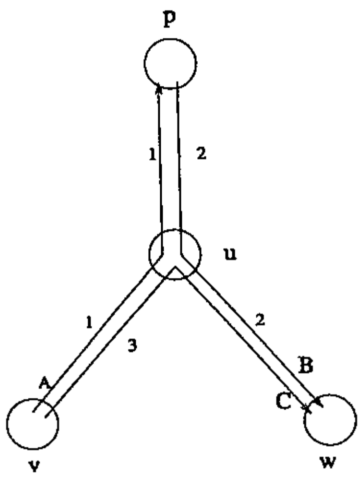

Fig. 4. A set of requests with load 2 that require 3 wavelengths.

Consider the wavelength converter $CONV_L = \{A_L, B_L, E_L\}$, with $|A_L| = |B_L| = L$ (Figure 5 depicts $CONV_9$). Partition the vertices of A_L and B_L in $\sqrt{L}$ subsets $A_1, A_2, \cdots, A_{\sqrt{L}}$ and $B_1, B_2, \cdots, B_{\sqrt{L}}$, where $A_i = \{a_{i1}, a_{i2}, \cdots, a_{i\sqrt{L}}\}$ and $B_i =$

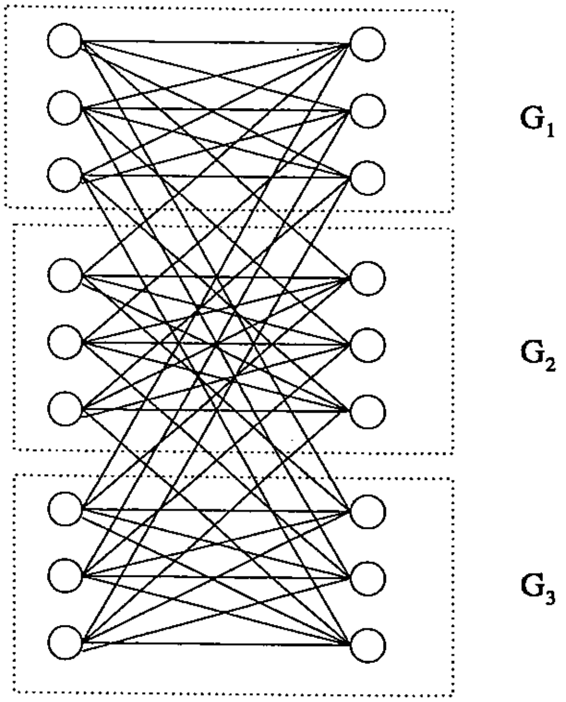

Fig. 5. The wavelength converter CONV designed for $L = 9$.

$\{b_{i1}, b_{i2}, \cdots, b_{i\sqrt{L}}\}$ for each $1 \leq i \leq \sqrt{L}$. Vertex a_{ij}, for $1 \leq i, j \leq \sqrt{L}$, is connected through an edge of E_L to vertex b_{il}, for each $1 \leq l \leq \sqrt{L}$, and to vertex b_{hj}, for $h \neq i$. It can be easily seen that the converter $CONV_L$ is a $(2\sqrt{L}-1)$-regular bipartite graph.

Lemma 3. *Assume that node u has a wavelength converter $CONV_L$ for each directed fiber link incident on it. Let A and B be the set of requests $v \to u$ and $p \to w$, respectively, and assume that each of these requests has been already colored. It is possible to convert the colors assigned to requests of A and B in such a way that they have at least* $\max\{|A|, |B|\}$ *common colors.*

The proof of the lemma is based on the next lemma. Consider the 3-level graph G_L obtained by putting two copies of $CONV_L$ back to back. More precisely, $G_L = (U_L, W_L, V_L, E_L)$, where $|U_L| = |V_L| = |W_L| = L$ and the two subgraphs induced by the vertices of U_L and W_L and by the vertices of W_L and V_L are isomorphic to $CONV_L$ and thus the vertices of U_L, W_L and V_L can be partitioned into blocks $U_1, \cdots U_{\sqrt{L}}, W_1, \cdots W_{\sqrt{L}}, V_1, V_2, \cdots V_{\sqrt{L}}$ in such a way that each vertex of U_j (V_j) is connected to all the vertices of W_j and to one vertex of each block W_h with $h \neq j$.

Lemma 4. *For each $l \leq L$ and each $X \subseteq U_L, Y \subseteq V_L$, with $|X| = |Y| = l$, there exist l vertex disjoint paths in G_L connecting vertices of X and Y.*

Proof. We prove the lemma by exhibiting an algorithm that, given sets of vertices X and Y of size l, constructs a set of l disjoint paths between X and Y. For each j, let $X_j = U_j \cap X$ and $Y_j = V_j \cap Y$ and denote by x_j the size of X_j and by y_j the size of Y_j. We first show how to find disjoint paths between the vertices of X and Y under the hypothesis that, for each j at most one between x_j and y_j is greater than 0. Then, we show how this construction can be generalized.

Consider the following path construction algorithm (without loss of generality we assume that $x_1 > 0$).

1. set h equal to 1 and k equal to 2
2. while $h \leq \sqrt{L}$ do
3.   while $x_h \geq y_k$ do
5.   connect y_k vertices of X_h to vertices of Y_k through W_k
6.   set x_h equal to $x_h - y_k$
7.   increment k
8.   connect the remaining vertices of X_h to x_h vertices of Y_k through W_h
9.   set y_k equal to $y_k - x_h$
10.   increment h

Notice that steps 5 and 8 are correct. In fact, each vertex of U_h is adjacent to a distinct vertex of W_k, while each vertex of V_k is adjacent to all the vertices of W_k. Thus it is always possible to find disjoint paths between any set of vertices of U_h and any set of vertices of V_k the same size that traverse vertices of W_k. Symmetrically, it can be proved that disjoint paths between X_h and Y_k exist through W_h. Notice that for each j the vertices of W_j are used to connect only vertices of X_h to vertices of Y_k for exactly one pair h, k such that j is equal to either h or k. Clearly these paths are disjoint from all the paths not containing vertices of W_h. Moreover, by the correctness of steps 5 and 8 it follows that they they are mutually disjoint.

Suppose now that x_j and y_j can be both greater than 0, for any j. Divide the set X_j in two subsets $X_{j'}$ and $X_{j''}$ and the set Y_j in $Y_{j'}$ and $Y_{j''}$, where $|X_{j'}| = |X_{j''}| = \min\{x_j, y_j\}$. We first use the previous algorithm to connect the vertices of $X_{1'}, X_{2'}, \cdots, X_{\sqrt{L}'}$ to the vertices of $Y_{1'}, Y_{2'}, \cdots Y_{\sqrt{L}'}$ (notice that for each j at least one between $X_{j'}$ and $Y_{j'}$ is equal to 0). Then, we connect vertices of $X_{j''}$ to vertices of $Y_{j''}$ through vertices of W_j not used in the previous step. Observe that these free vertices exist, since at most $y_j - \min\{x_j, y_j\}$ vertices of W_j have been used in the first step of the assignment.

Proof of Lemma 3: Consider the graph G_L and let X be the set of vertices of U_L corresponding to the colors assigned to the requests of A, and let Y be the set of vertices of V_L corresponding to the colors assigned to the requests of B. By

Lemma 4 there are disjoint paths connecting vertices of X to vertices of Y. Thus, for each $x \in X$ there is a request of A with color x and a request of B with color $y \in Y$ such that one of the paths obtained in Lemma 4 connects x to y through $w \in W$. Then, we can convert color x of A and color y of B to color w. □

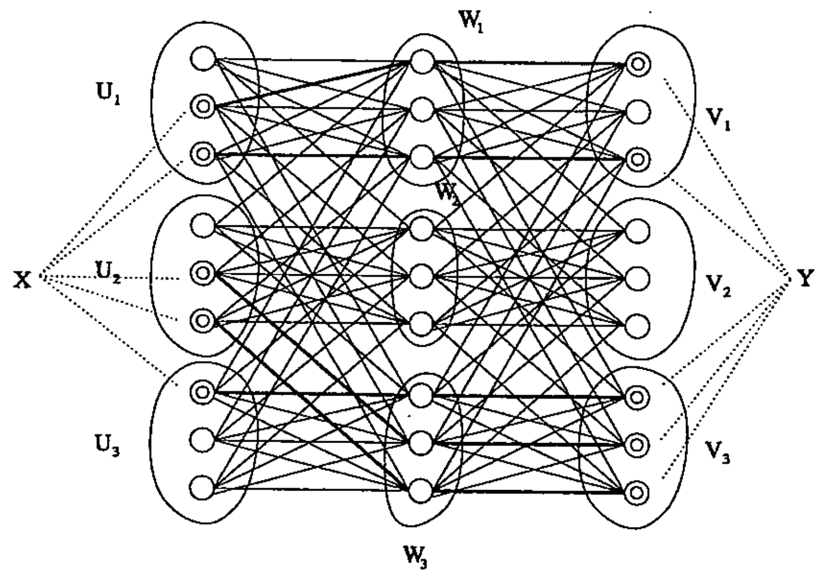

Fig. 6. The graph G_9 obtained by putting two copies of $CONV_9$ back to back.

Finally, we can state the main result of this section as follows.

Theorem 6. *Any set of requests of load L on a binary tree T can be greedily routed with L wavelengths using converters of degree* $2\sqrt{L}-1$.

Proof. Let v be a vertex of T and let p,w,u the vertices adjacent to v and let $R_1,R_2,\cdots,R_6$ be sets of requests that touch vertex u, as described in Figure 7. We prove the theorem by showing that the algorithm can color all the requests in $R_1,\cdots R_6$ using only L colors.

By construction when the algorithm visits vertex u requests of R_1,R_2,R_3, and R_4 have already been colored. By Lemma 3, it is possible to convert the colors used for the requests of R_1 and R_4 in such a way that they have at least $\min\{|R_1|,|R_4|\}$ in common. Therefore, there are at least $L-\max\{|R_1|,|R_4|\}$ colors that are available both on the fiber link (v,u) and on (u,w). But, since the load of R is L, we have that $|R_5| \le L-\max\{|R_1|,|R_4|\}$ and thus the algorithm can assign an available color to all requests of R_5.

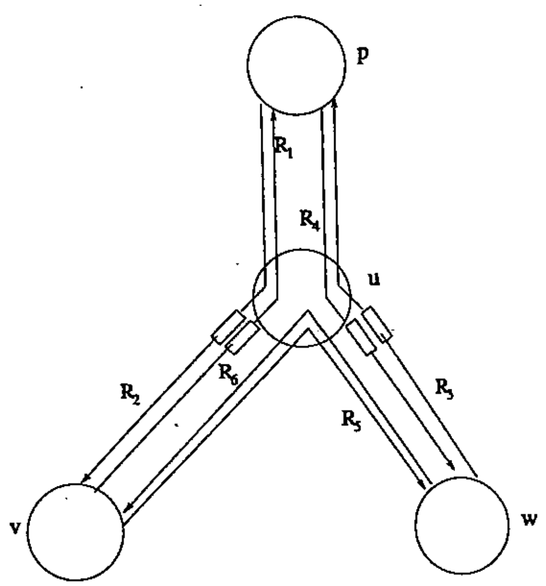

Fig. 7. A node u along with the requests that touch u and the converters on the edge that are adjacent to u.

Similarly, we can prove that it is possible to convert the colors assigned to R_2 and R_3 in such a way that there is a number of available colors on the fiber links (w, u) and (u, v) that is greater than the size of R_6. Thus, it is possible to assign an available color to each request of R_6 while using at most L colors.

References

1. A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswani, B. Shieber, and M. Sudan, *Efficient Routing and Scheduling Algorithms for Optical Networks.* In Proc. of SODA 93.
2. J.C. Bermond, L. Gargano, S. Perennes, A.A. Rescigno, and U. Vaccaro, *Efficient Collective Communication in Optical Networks.* In Proc. of ICALP 96.
3. I. Caragiannis, C. Kaklamanis, P. Persiano, *Bounds on Optical Bandwidth Allocation on Directed Fiber Tree Topologies.* In Proc. of the 2nd Workshop on Optics and Computer Science, IPPS 1997.
4. K.W. Cheung et al., *IEEE JSAC: Special Issue on Dense WDM Networks.* Vol. 8, 1990.
5. G.K. Chung et al., ONTC-ARPA Technology Development Group, *Experimental Demonstration of a Reconfigurable WDM/ATM/SONET Multiwavelength Network Testbed.* OFC Optical Fiber Conference, San Hose CA, January 1993.
6. P.E. Green, *Fiber Optics Communication Networks.* Prentice Hall, 1992.
7. K. Jansen in Proc. of the 2nd Workshop on Optics and Computer Science, IPPS 1997.

8. C. Kaklamanis, P. Persiano, *Efficient Wavelength Routing on Directed Fiber Trees.* In Proc. of ESA 1996.
9. C. Kaklamanis, P. Persiano, T. Erlebach, K. Jansen, *Constrained Bipartite Edge Coloring with Applications to Wavelength Routing in All-Optical Networks.* In Proc. of ICALP 1997.
10. L.G. Kazovski, P.T. Poggiolini, *STARNET: A Multi–Gigabit–per–second Optical LAN utilizing a Passive WDM Star.* J. of Lightwave Technology, 11, pp. 1009–1027, 1993.
11. V. Kumar, E. Schwabe, *Improved Access to Optical Bandwidth in Trees.* In Proc. of SODA 97.
12. K. -C. Lee, V.O.K. Li, *Routing and Switching in a wavelength convertible lightwave network.* In Proc. of IEEE Infocom 1993, pp. 578–585.
13. M. Mihail, C. Kaklamanis, S. Rao, Efficient Access to Optical Bandwidth. In Proc. of FOCS 1995.
14. D. Minoli, *Telecommunications Technology Handbook.* Artech House, 1991.
15. R. Pankaj, *Architectures for Linear Lightwave Networks.* Ph.D. Thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge MA, 1992.
16. P. Raghavan, E. Upfal, *Efficient Routing in All-Optical Networks.* In Proc. of STOC 1994.
17. R. Ramaswami, G.H. Sasaki, *Multiwavelength Optical Networks with Limited Wavelength Conversion.* In Proc. of IEEE Infocom 1997.
18. S. Subramaniam, M. Azizoglu, A.K. Somani, *Connectivity and Sparse Wavelength Conversion in Wavelength-Routing Networks.* In Proc. of IEEE Infocom 1996, pp. 148-155.
19. J. Yates, J. Lacey, D. Everitt, M. Summerfield, *Limited-Range Wavelength Translation in All-Optical Networks.* In Proc. of IEEE Infocom 1996, pp. 954-961.
20. J. Zhou, N. Park, K. Vahala, M. Newkirk, B. Miller, *Four–Wave Mixing Wavelength Conversion Efficiency in Semiconductor Travelling–Wave Amplifiers Measured to 65 nm of Wavelength Shift.* IEEE Phot. Tech. Letters, vol. 6, pp. 984–987, August 1994.