A New Model for Optimal Routing and Wavelength Assignment in Wavelength Division Multiplexed Optical Networks[1]

Chien Chen and Subrata Banerjee

Department of Electrical Engineering and Computer Science
Stevens Institute of Technology, Hoboken, NJ 07030

Abstract

We consider the problem of routing and assignment of wavelength (RAW) in optical networks. Given a set of requests for all–optical connections (or lightpaths), the problem is to (a) find routes from the source nodes to their respective destination nodes, and (b) assign wavelengths to these routes. Since the number of wavelengths is limited, lightpaths cannot be established between every pair of access nodes. In this paper, we first consider the dynamic RAW problem where lightpath requests arrive randomly with exponentially distributed call holding times. Then, the static RAW problem is considered which assumes that all the lightpaths that are to be set–up in the network are known initially. Several heuristic algorithms have already been proposed for establishing a maximum number of lightpaths out of a given set of requests. However, most of these algorithms are based on the traditional model of circuit–switched networks where routing and wavelength assignment steps are decoupled. In this paper a new graph–theoretic formulation of the RAW problem, dubbed as layered–graph, has been proposed which provides an efficient tool for solving dynamic as well as static RAW problems. The layered–graph model also provides a framework for obtaining exact optimal solution for the number of requested lightpaths as well as for the throughput that a given network can support. A dynamic and two static RAW schemes are proposed which are based on the layered–graph model. Layered–graph–based RAW schemes are shown to perform better than the existing ones.

1. Introduction

All–optical local and wide area network infrastructures supporting hundreds and thousands of users, each operating at gigabit–per–second speed, can be realized by wavelength–division–multiplexing (WDM) and other emerging optical communications technologies[1]. In this paper, we consider a WDM optical network architecture consisting of wavelength routers which are interconnected by point–to–point optical links (*e.g.*, see Fig. 1). End users or electronic switching devices (*e.g.*, ATM switches) can be attached to the wavelength routers via *access* nodes. Direct optical connection between a pair of access nodes can be set–up in this all–optical network by appropriately choosing a wavelength–continuous route. Since no optical–to–electrical conversions or buffering is performed at any intermediate wavelength routing node, a Terabits–per–second network supporting a large number of users at gigabit–per–second rate can be achieved. For example, in Fig. 1, an optical path (also called *lightpath* or *clear channel*) can be set–up between Node *B* and Node *C* without any intermediate electronic processing.

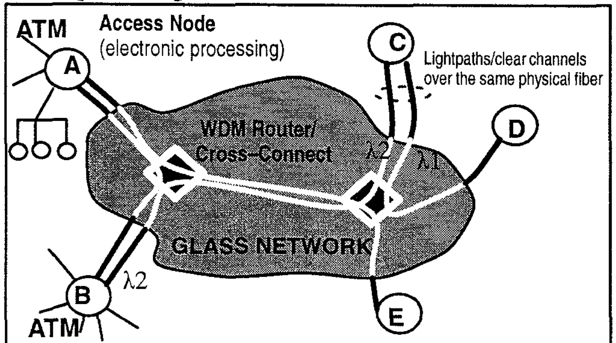

Fig. 1. An example architecture of all–optical wavelength routing network.

Since the number of wavelengths is limited, it may not be possible to set up one–hop[2] connection between every pair of access nodes. As a result, multiple hops[2] may be necessary for some of the communicating access node pairs. Therefore, based on some criterion (*e.g.*, network traffic demands[3]), a set of lightpaths has to be selected for embedding a virtual topology on top of a physical topology such that the network performance (e.g., throughput, blocking rate, average message delay etc.) is optimized. Thus, the goal for *a dynamic RAW* scheme is to minimize the blocking probability, and one of the goals for a *static RAW* scheme is to maximize the total one–hop traffic for a given traffic demand matrix.

Routing and ***A***ssignment of ***W***avelength (*RAW*) problem in WDM optical networks can be described as follows. Given a set of requests for all–optical connections or lightpaths between access node pairs, the problem is to (*a*) find routes from the source nodes to their respective destination nodes, and (*b*) assign wavelengths to these routes. In this paper, we first consider a *dynamic RAW* problem in which lightpath requests arrive randomly with random call holding times. Then, a *static RAW* problem is considered which assumes that all the lightpaths that are to be set–up in the network are known initially.

A similar routing problem is typical in traditional circuit–switched telephone networks. Consider an undirected *multigraph* model of a physical network, where multiple

1. This work has been partially supported by NSF under Grant No. NCR–9496268.

0743-166X/96 $5.00 © 1996 IEEE

edges between same pairs of nodes correspond to multiple logical channels (within the same physical cable). A call between two users can be established by finding a route in the multigraph that is edge–disjoint with the existing connections. However, in an optical network, the path, apart from being edge–disjoint with the existing connections, must also satisfy the wavelength continuity constraint, which demands that the same wavelength should be assigned to all the segments of the path. For example, in Fig 2(a), let us assume that each physical link in a traditional circuit–switched network supports two logical channels. In the corresponding undirected multigraph (Fig. 2(b)), three connections can be accommodated by finding three edge–disjoint paths. However, in optical networks, the paths must use the same wavelength on all the links. With this additional constraint, only two of the three edge–disjoint paths can be accommodated with two wavelengths (*i.e.*, optical channels) per link.

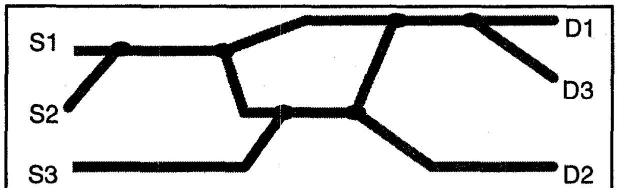

Fig 2(a). Physical Network. Each line represents an optical fiber, each fiber can support two WDM channels.

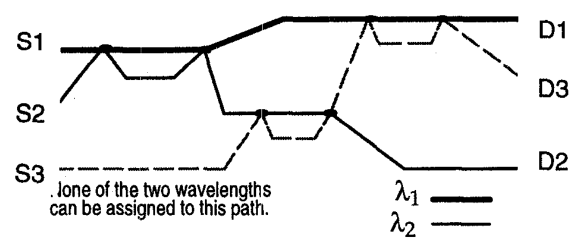

Fig 2(b). Logical Network. Each line represents a WDM channel. Although there are three edge disjoint paths, only two of them can be assigned as lightpaths.

Several researchers have studied the *RAW* problem and various heuristic algorithms have been developed [3,4,5,6]. Most of these algorithms are based on the undirected multigraph model of traditional circuit–switched networks. In this model, the routing and wavelength assignment steps are decoupled. That is, first the best route (*e.g.*, the shortest path) for each lightpath request is found, and then wavelengths to these paths are assigned one–by–one based on some heuristic criterion such as 'longest–path first' [6] or 'heaviest traffic path first' [3]. In [5], *W* copies of the physical network topology are created, one for each of the *W* available wavelengths. Thus, the optical network is represented as *W* disjoint subgraphs, each of which is identical to the physical network topology. A route between a pair of nodes is found by applying a shortest path algorithm on all the *W* disjoint subgraphs, and then the cheapest one among these shortest paths is chosen. However, this disjoint subgraph model works only in the case of dynamic *RAW* problem, where the lightpaths are assigned on demand. This model is not suitable for optimal static *RAW* problem, since each subgraph is considered independently and there is no way to decide which lightpath request should be assigned to which subgraph in order to accommodate a maximum number of lightpaths in the optical network.

Given that the number of wavelengths is limited, it is unlikely that lightpaths can be established between every pair of access nodes. Moreover, a *common* wavelength might not be available on all the segments of a chosen route (a *segment* is a portion of the route between two adjacent wavelength routers). Thus, routing and wavelength assignments are to be considered simultaneously for the best possible performance. In this paper, we propose a new model called *layered–graph* for addressing the dynamic as well as static *RAW* problems. We show that under this new model, the *RAW* problem can be reduced to the problem of finding edge–disjoint paths in a traditional circuit–switched network. Thus, the routing algorithms used in the traditional circuit–switched networks can also be applied to our layered–graph model. Based on the *layered–graph* model, a maximal multicommodity 0–1 flow formulation of the static *RAW* problem has also been obtained. An exact optimal solution for a given *RAW* problem can be obtained by solving the corresponding integer linear programming problem. To our knowledge no other formulation exists that can provide an exact optimal solution to the *RAW* problem.

The remainder of this paper is organized as follows. The *layered–graph* model is introduced in Section 2. Section 3 deals with the *dynamic RAW* problem. The *static RAW* problem is considered in Section 4. Numerical results for *dynamic* as well as *static RAW* schemes are presented in Section 5. The paper concludes in Section 6.

2. A New Model For Solving *RAW* Problem

In the last section, we observed that solving the optimal *RAW* problem is quite difficult due to its wavelength continuity constraint. In fact, it has been proved that the optimal *RAW* problem is NP–complete[6]. Several schemes have already been proposed for establishing a maximum number of desired lightpaths in a given topology. These schemes first find routes for a given set of lightpath requests, and then the schemes try to establish as many lightpaths as possible along the *predetermined* routes. Thus, the routing and the wavelength assignment subproblems are *decoupled.* Note that, even if the maximum number of paths routed via any link in a given topology is *W*, it is not guaranteed that all these routes can be supported with *W* wavelengths (see Fig. 2). In the *layered–graph* model a lightpath request is routed based on the available links as well as on the available wavelengths in these links such that, whenever a route is found its wavelength continuity is guaranteed. Thus, in the *layered–graph* model, routing and wavelength assignment steps are tightly coupled.

The *layered–graph* model also provides a framework for obtaining *exact upper bound* on the number of requested lightpaths a given topology can support. Another formulation for this upper bound can be found in [4] which, however, is applicable only for large number of wavelengths.

2.1. The Layered–Graph Model

Define a network topology *N(R, A, L, W)* for a given WDM optical network (e.g., see Fig. 1),

where **R** is the set of wavelength router nodes,
A is the set of access nodes,
L is the set of undirected links, and
W is the set of available wavelengths per link.

At a wavelength router $r \in \boldsymbol{R}$, wavelengths arriving at its different input ports can be routed onward to any of its output port, provided that all the wavelengths routed to the same output port are distinct. Each access node in **A** is attached to a wavelength router. The access nodes provide electro–optical conversions for supporting *electronic* packet and circuit switching operations (e.g., electronic cell switching in ATM/BISDN). The links, **L**, in the network are assumed to be bidirectional. Each link consists of two unidirectional optical fibers, each carrying $|\boldsymbol{W}|$ wavelength division multiplexed channels.

The ***layered–graph*** model is a *directed* graph, ***G(V, E)***, obtained from a given network topology *N* as follows. Each node $i \in \boldsymbol{R}$ in *N*, is replicated $|\boldsymbol{W}|$ times in **G**. These nodes are denoted as $v_i^1, v_i^2, .. v_i^W \in V$. If link $l \in \boldsymbol{L}$ connects Router *i* to Router *j*, $i, j \in \boldsymbol{R}$, in *N*, then Node v_i^w is connected *to* Node v_j^w by a directed edge, e_{ij}^w, $e_{ji}^w \in \boldsymbol{E}$ for all $w \in W$. Now, let us consider an access node $a \in A$ attached to Router *k*. In the layered–graph **G**, two nodes are created for each access node *a*, one representing the traffic generating part (*i.e.,* source) of Node *a*, while the other represents the traffic absorbing part (*i.e.,* destination). These two nodes are denoted as v_a^s and $v_a^d \in V$, respectively. Now, *directed* edges from Node v_a^s to nodes $v_k^1, v_k^2, .. v_k^W \in V$, and from each of the nodes $v_k^1, v_k^2, .. v_k^W \in V$ to v_a^d are added to **G**. Thus, the number of nodes in **G** is $|V| = |\boldsymbol{R}| \times |\boldsymbol{W}| + 2|\boldsymbol{A}|$ and the number of directed edges in **G** is $|\boldsymbol{E}| = 2|\boldsymbol{L}| \times |\boldsymbol{W}|$. Fig. 3(b), shows an example of a layered–graph which is obtained from the physical network topology shown in Fig. 3(a).

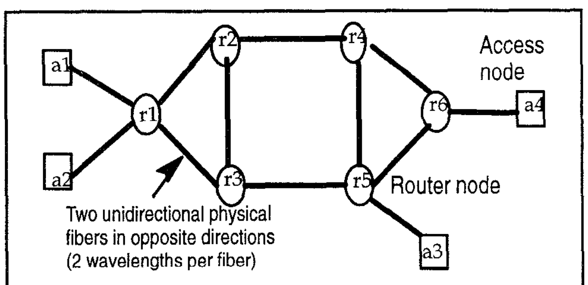

Fig. 3.(a) A given optical network topology *N*.

It is easy to see that if a set of paths in the layered–graph is edge–disjoint, then all the paths in the set can be supported in the corresponding physical network topology. However, recall that in the multigraph model, a set of edge–disjoint paths does not guarantee that they all can be supported in the *physical* network topology, *N* (*e.g.,* see Fig. 2).

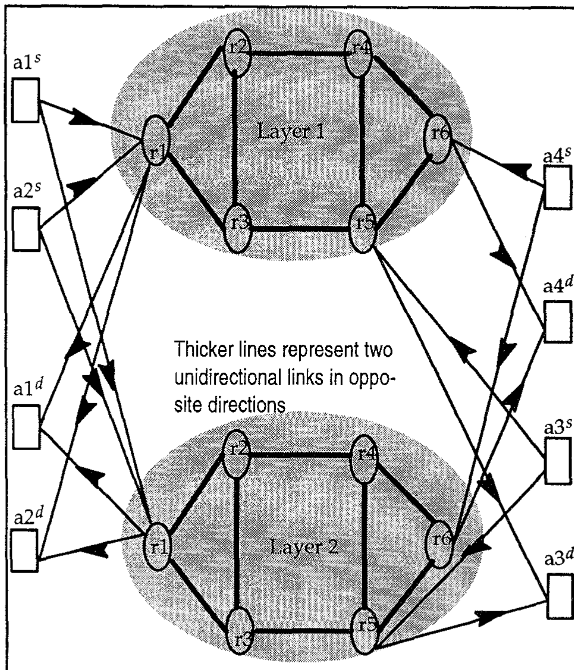

Fig. 3.(b) Layered–graph corresponding to the given optical network topology *N*.

Before we proceed further we need to introduce the following terminologies.

Definition: A **Mono–Chromatic Path (MCP)**, from access Node *i* to access Node *j* in *N*, is a directed path from Node *i* to Node *j* such that all the links in the path can be assigned a common wavelength (color) without conflicting with any existing MCPs in the network.

Let $\boldsymbol{M} = \{(s_i, d_i) \mid s_i, d_i \in \boldsymbol{A}\}$ be a set of lightpath requests from access node s_i to access node d_i. (Note that multiple lightpaths can be requested between two access nodes.) A necessary and sufficient condition for the existence of MCPs for a given set of lightpath requests *M* is given by the following theorem. A simple proof of the theorem is given in the appendix.

***Theorem* 1**: Let **G** be a layered–graph constructed from an optical network topology *N* and let **M** be a set of lightpath requests. Then, *N* can support all the lightpath requests in **M** as MCPs if and only if **G** contains a set of edge–disjoint paths for all the requests in **M**.

After introducing the necessary terminologies we now consider the dynamic RAW problem in Section 3 and the static RAW problem in Section 4. Numerical results are presented in Section 5.

3. The Dynamic *RAW* Problem

In the dynamic RAW problem lightpaths between access nodes are set up on demand. Lightpath requests are assumed to arrive according to a Poisson process with exponentially distributed holding times. Whenever a new lightpath is requested the dynamic *RAW* scheme either selects an available MCP, or it blocks the request if no such path can be found. However, if multiple MCPs exist, a *dynamic RAW* scheme should look for the "best" MCP (*e.g.,* the one that is likely to minimize blocking in the future). Unfortunately, finding an MCP in the network, *N,* is not straightforward. First, a best route (*e.g.,* the shortest one) has to be found, and then one has to decide if any wavelength can be assigned to that route. If none of the wavelengths could be assigned to that route due to the wavelength continuity constraint, then the second best route has to be found, and so on. These procedures might have to be repeated several times until either an MCP is found, or the request is blocked. Searching time for an MCP in this manner can be large especially when the network already contains a large number of MCPs.

The dynamic *RAW* problem can be simplified significantly by employing the layered–graph model. We show that a set of edge–disjoint paths in a layered–graph guarantees monochromaticity of the paths in the corresponding optical network. In other words, if a route, that is edge–disjoint with the existing connections, cannot be found in the layered–graph, then that connection request cannot be met. Thus, the *dynamic RAW* problem can be solved by simply applying a shortest path algorithm for the requested lightpath in the *layered–graph.* A formal description of this formulation is provided in Subsection 3.1. The layered–graph–based *dynamic RAW* scheme is presented in Subsection 3.2.

3.1. Mathematical formulation

Consider an optical network topology ***N(R, A, L, W)*** with cost *c(l)* associated with each of its optical link $l \in L$. Let ***B*** be the set of lightpaths that are already set–up in ***N***. Also, at time *t*, let $M = \{m=(s_i,d_i),\ s_i,d_i \in A\}$ be a set of source–destination pairs requesting new *lightpath* connections. Let ***G(V, E)*** be the layered–graph constructed from ***N***. Now, weights for the edges in ***G*** are assigned as follows. If an existing lightpath is routed via wavelength *k* on link $l = (i, j)$ in ***N***, then the weight of the directed edge $e^k_{ij} \in E$ in ***G*** is set to infinity, otherwise, the weight of e^k_{ij} is set to *c(l).*

Now, the *dynamic RAW* problem can be formulated as follows. For each node pair *(s,d)* in ***M***, our goal is to set up the 'cheapest' *available* lightpath between Node *s* and *d*. Alternatively, the problem can be viewed as sending one unit of flow from Node *s* to Node *d* via the least cost path in the layered–graph. Let x^k_{ij} be defined as follows.

$$x^k_{ij} = \begin{cases} 1 & \text{if 1 unit of flow passes through the edge } e^k_{ij} \\ 0 & \text{otherwise} \end{cases}$$

Then, the minimal cost, C_p, of a lightpath from Node *s* to Node *d* can be formulated as follows:

$$\begin{aligned} C_p &= \min \sum_{i,j \in A,\ k \in W} c(e^k_{ij}) \cdot x^k_{ij} \quad \text{subject to} \\ \sum_{j \in V} x^k_{ij} - \sum_{j \in V} x^k_{ji} &= \begin{cases} 1 & \text{if } i = s \\ -1 & \text{if } i = d \\ 0 & \text{otherwise} \end{cases} \\ x^k_{ij} &\geq 0 \quad \text{for all } i,j \end{aligned} \qquad (1)$$

Note that, since the links to the *access* nodes in ***V*** are unidirectional, a lightpath from access node *s* to access node *d* cannot be routed via another access node. This enforces the wavelength contuinity constraint since it is guaranteed that every route will be entirely in one of the layers in ***G***. This was the reason for creating two nodes in the layered–graph ***G***, one as the source and the other as the sink, for every access node in ***N***.

3.2. The dynamic *RAW* algorithm

The above formulation can be solved by a shortest path algorithm (*e.g.,* Dijkstra's algorithm). If the cost of the shortest path, C_p, is found to be finite then the request is accepted and the lightpath is set–up along the shortest path; otherwise the request is blocked. Note that, the weights on the links of *G* have to be updated whenever a lightpath is established or released. A formal description of the algorithm is given below:

Layered–Graph based *dynamic RAW* algorithm:

1. Transform a given WDM optical network ***N(R, A, L W)*** into a layered–graph ***G(V, E).***
Let the cost[2] of edge *e* in ***G*** be *c(e),* for all $e \in E$
2. Wait for a lightpath request.
If it is a lightpath connection request, go to Step 3.
If it is a lightpath release request go to Step 4.
3. Find a shortest path *p* in ***G*** from the source to the destination node (*e.g.,* by Dijkstra's algorithm).
If the cost of the path $C_p = \infty$, reject the request; otherwise, Accept the request
  Set up the connection along the shortest path.
  If the shortest paths is via Layer *k*, assign wavelength *k* to the lightpath
  Update the cost of the edges on path *p* to ∞.
Go to Step 2.
4. Update the cost of the edges e_i occupied by the lightpath to $c(e_i)$.
Release the lightpath and then go to step 2.

The **complexity** of step 1 is linear with the number of edges |**E**| in ***G***. Also, the running time of Step 4 is at most proportional to the diameter of the network *N*. Thus, the computational complexity of the *dynamic RAW* algorithm is dominated by the point–to–point shortest path algorithm used in Step 3. If $c(e_i) = 1$ for all *i*, then Step 3 can be executed in O(|***E***| +|***V***|) time using Dial's implementation of Dijkstra's algorithm. If the lengths of the physical links are used as the cost function then Step 3 can be executed in O(|***E***|+|***V***|***D***) where ***D*** is the diameter of ***G***[7]. Note that, in

2. This cost may be based on the physical length of the link if propagation delay is a concern or the cost can be set to 1 for all $e \in E$ if the number of wavelength routers in a lightpath is to be minimized.

large or dense optical networks the total number of links is much greater than the number of wavelengths per link. In such cases the complexity of the layered–graph–based algorithm is within a constant factor of the complexity of routing algorithms for traditional circuit–switched networks.

4. The Static *RAW* Problem

In *static RAW* problem it is assumed that set of the requested lightpaths, M, to be set–up in the network is known initially. The objective here is to maximize network throughput. We consider two traffic cases, *uniform* and *non–uniform.* In *uniform traffic* case each lightpath has equal traffic demand. Thus, maximizing network throughput is same as maximizing the number of lightpaths established. In the *non–uniform traffic* case traffic demands of lightpath requests are randomly distributed. Here our objective function is to support as much as total carried traffic as possible. Based on the layered–graph model, static *RAW* problem can be formulated as a maximal multicommodity 0–1 flow problem. Note that, a similar integer linear programming formulation can be found in [4]. The formulation in [4] is based on the number of maximal independent sets in a path–graph[3][8] which can grow exponentially with the number of nodes in path–graph. A formal description of our layered–graph–based formulation of the static *RAW* problem is given in Subsection 4.1. Two layered–graph–based static *RAW* schemes are presented in Subsection 4.2 These two schemes are evaluated and compared against the upper bound in [9].

4.1. Mathematical formulation

Consider an optical network $N(R, A, L, W)$, and a set of source–destination access node pairs M. Based on the layered–graph model, the static *RAW* problem can be reduced to the problem of finding a maximal set of edge–disjoint paths in G, connecting the node pairs in M. A *multicommodity* 0–1 *flow* based formulation of the *static RAW* problem is given below.

Let $G(V, E)$ be a layered–graph corresponding to a given physical topology N. Let the capacity of the edges in E be unity. Also let $f(s_n, d_n)$, $n = 1,\ldots,|M|$, denote the value of flow n carried from access node s_n to access node d_n where,

$$f(s_n, d_n) = \begin{cases} 1 & \text{if a lightpath is assigned from } s_n \text{ to } d_n \\ 0 & \text{otherwise} \end{cases}$$

Also, let x^n_{ij} denote the value of flow n on the edge $e_{ij} \in E$ where,

$$x^n_{ij} = \begin{cases} 1 & \text{if flow n is assigned to edge } e_{ij} \\ 0 & \text{otherwise} \end{cases}$$

3. In a path graph, $G_p = (V_p, E_p)$, each node corresponds to a lightpath in the optical network and an edge connects two nodes in G_p when the two corresponding lightpaths share a common physical link.

Now, the optimal static *RAW* problem for uniform traffic case can be formulated as follows.

$$\text{maximize} \sum_{n=1}^{|M|} f(s_n, d_n) \quad \text{subject to}$$

$$\sum_{i \in V} x^n_{ij} - \sum_{k \in V} x^n_{jk} = \begin{cases} f(s_n, d_n) & \text{if } j = d_n \\ -f(s_n, d_n) & \text{if } j = s_n \\ 0 & \text{otherwise} \end{cases}$$

$$\sum_{n=1}^{|M|} x^n_{ij} \leq 1 \quad \text{for all } i, j$$

$$x^n_{ij} \geq 0 \quad \text{for all } i, j \qquad (2)$$

For non–uniform traffic model, traffic demand between each source–destination access node pair is given in the form of a traffic demand matrix $T = [t_{ij}]$. We also assume that the traffic demand between each pair of nodes is less than the bandwidth which a lightpath provides. Otherwise, we will split the request to several requests between the same source–destination node pair. Let $t(s_n, d_n)$, $n = 1,\ldots,|M|$, denote the traffic demand from an access node s_n to an access node d_n. The optimal solution can be formulated as follows:

$$\text{maximize} \sum_{n=1}^{|M|} t(s_n, d_n) f(s_n, d_n) \quad \text{subject to (2)}$$

Raghavan and Upfal[10] conjectured that the static *RAW* problem is significantly harder than integer multicommodity flow problem due to the wavelength continuity constraint. However, our formulation of the optimal static *RAW* problem is a pure integer multicommodity flow problem with respect to the layered–graph model. The above formulation can be solved as an integer linear programming (ILP) problem. Thus, an upper bound on the throughput supported by the network can be obtained by relaxing the ILP to a linear program.

4.2. The static *RAW* algorithms

For large networks the running time of the ILP is likely to be unacceptably long. In such cases, heuristic or approximate algorithms for integer multicommodity flow problem can be used to obtain faster but suboptimal solution. Study of algorithms for maximal multicommodity 0–1 flow is outside the scope of this paper. However based on the layered–graph, two simple heuristic algorithms, one for the uniform traffic model and the other for non–uniform traffic model, are proposed. These algorithms have been simulated in order to better understand the advantages of the layered–graph model in solving the static *RAW* problem.

For uniform traffic case, a hybrid algorithm which combines greedy and layered–graph approach is proposed. This algorithm, in its first phase, finds shortest paths for all lightpath requests in given network topology N. Then the lightpath requests are sorted in non–descending order of their shortest path lengths. The algorithm now routes first lightpath request on the layered–graph with the shortest path. If a finite–cost path is found then the lightpath is established, and the layered–graph is updated by assigning infinite cost to the edges along which the request is routed. Otherwise the request is skipped and the next request from the sorted list is considered. The procedure is repeated until all the lightpath requests are considered. In the second phase, the algorithm searches the shortest *available* path in the *residual* layered–graph for the lightpath requests that were skipped in the first phase. The details of the steps of the algorithm are described in [9].

Next we consider non–uniform traffic case. Here our objective is to maximize the total carried traffic in the network. Therefore, we first order the lightpath requests by their traffic demands. Our layered–graph based heuristic algorithm then tries to satisfy the lightpath requests in the non–decreasing order of their traffic demands. Detailed description of algorithm and its numerical performance can be found in [9].

5. Numerical Results

In this section numerical results of the layered–graph based *dynamic* and *static RAW* schemes are presented. An ARPANET–like network topology is used as an example[3]. This network, shown in Fig. 4, consists of 24 wavelength routers interconnected by 49 bidirectional optical links. We also assume that six access nodes are attached to each wavelength router (to keep the figure simple these access nodes are not shown in Fig. 4).

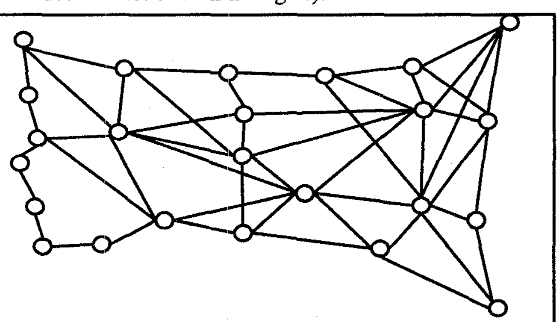

Fig. 4. A 24–node ARPANET–like network topology *N*.

For the dynamic case, lightpath requests are assumed to arrive according to an independent Poisson process with arrival rate *r*. Holding time of the lightpaths are exponentially distributed with mean $\frac{1}{\mu}$ (in the simulations average holding time is set to 1 time unit). Also, once a lightpath request is blocked, it is removed from the list of pending lightpath requests. The traffic pattern is assumed to be uniform, that is, the lightpath requests are uniformly distributed over all the access node pairs.

Performance of the layered–graph based *dynamic RAW* scheme is shown is Fig. 5. Performance of the simple greedy heuristic[4], which accepts a lightpath request if a common wavelength is available on all the segments of a shortest path is also shown. Fig. 5(a) shows the blocking rate of the lightpath requests as a function of their arrival rate *r*. Fig. 5(b) shows the average number of ongoing lightpath connections as a function of their arrival rate. As expected, the results show that the layered–graph–based approach has a better chance in finding an MCP for a requested lightpath. Thus, for the same arrival rate, more lightpath requests can be satisfied by using the layered–graph–based approach. For example, when $r = 100$ and $|W| = 6$, the blocking rate of the layered–graph based algorithm is only 0.09 compared to blocking rate of 0.24 for the simple heuristic in [4]. Thus, by using the layered–graph based approach, blocking rate is reduced by 62%. Moreover, relative performance of the layered–graph based approach improves as $|W|$ increases.

Figs. 6 and 7 show the performance of the layered–graph–based *static RAW* heuristics for the uniform and non–uniform cases. The results are compared with two greedy approaches. Both the greedy approaches first find shortest paths for the all lightpath requests in the original network topology. Then, the first greedy heuristic assigns the wavelengths to lightpath requests in increasing order of their shortest distances[4], while the second heuristic assigns wavelengths in decreasing order of their traffic demands[3]. All the heuristic algorithms are simulated for various sets of source–destination access node pairs, $\boldsymbol{M}$, in the 24–node ARPANET–like example network topology. During the simulations, first a set of lightpath requests is generated randomly and uniformly selecting access node pairs. For the non–uniform traffic case, traffic demand for each of the requested lightpaths is set to be uniformly distributed between 0 and 1.

We plot the number of lightpaths established versus the number of lightpaths requested in Fig. 6(a) for the uniform traffic case. The four curves in Fig. 6(a) represent two different heuristic schemes (*i.e.* layered–graph and greedy[4] approaches) for $|W| = 6$ and 12. The simulation results show that the improvement obtained from the layered–graph–based approach, in term of the number of lightpaths established, is very consistent after a certain size of the lightpath request set. The average gain in the number of lightpaths established is about 22.2 for $|W| = 6$ and $160 \leq |M| \leq 840$. For $|W| = 12$ and $320 \leq |M| \leq 1000$, the layer–graph–based approach establishes on an average, 37.6 more lightpaths, which is a little less than twice the gain for $|W| = 6$. For example, for $|M| = 800$, the layered–graph–based algorithm can support 22 and 41 more lightpaths than the greedy algorithm when $|W| = 6$ and 12, respectively.

Fig. 6(b) shows the average hop distance (in terms of the number of physical links) of the established lightpaths versus the total number of lightpaths requested. The average hop distance for the layered–graph–based algorithm is expected to be longer than the greedy algorithm. (Recall that the layered–graph–based approach looks for an available monochromatic path which might be longer than the shortest path that the greedy algorithm insists on.) However, the simulations show that the largest differences in the hop distance are only 0.425 and 0.312 hop for $|W| = 6$ and 12, respectively; and the difference gets smaller as the number of lightpath request grows. Since the layered–graph–based algorithm also searches for an *available* shortest path in the corresponding layered–graph, increase in the average hop distance is mostly contributed by a few 'long' lightpaths which could not be satisfied by the greedy approach.

For the non–uniform traffic case, Fig. 7(a) plots the total traffic demands of the established lightpaths versus the number of lightpaths (*i.e.*, throughput) requested (*i.e.*, $|M|$). We find that the throughput of the layered–graph based approach is always greater than the throughput of the greedy algorithm[3]. This increased throughput is not only due to a greater number of lightpaths being supported but also due to the fact that more lightpaths with heavy traffic demands were established. From Fig. 7(b), which shows the total number of lightpaths established versus the number of lightpaths requested, we note that the improvement in the total number of lightpaths established decreases as the number of lightpath requests increases. However, the improvement in total network throughput shown in Fig. 7(a) is still significant and it remains constant for large values of $|M|$.

6. Conclusions

We presented a new layered–graph model for solving the routing and wavelength assignment problem in optical networks. The layered–graph model 'eliminates' the wavelength continuity constraint since all the paths between two access nodes are guaranteed to be monochromatic. Thus, the optimal *RAW* problem can be reduced to the problem of finding maximal number of edge–disjoint paths in the corresponding layered–graph. This allows us to consider the routing and wavelength assignment subproblems simultaneously in order to exploit better performance. This is in contrast to the traditional schemes in which the routing and wavelength assignment subproblems are treated independently. Based on the layered–graph model, the *RAW* problem is formulated as a pure maximal 0–1 multicommodity flow formulation. Heuristic algorithms that are based on the layered–graph model are also developed for the *dynamic* and *static RAW* problems. Since the number of wavelength is much smaller than the number of nodes in the network, the complexity of the layered–graph–based approach is within a small factor of the complexity of traditional schemes. Thus, the cost of routing in the layered–graph is only a few times the cost of routing in the physical network. Numerical examples were employed for evaluating the performance of the layered–graph–based heuristics with respect to the upper bounds and other existing heuristics. The simulation results demonstrate that the layered–graph–based algorithms provide better performance, both in terms of blocking rate and the number of lightpaths established (for the dynamic case) as well as in terms of the one–hop traffic supported (for the static case).

Acknowledgment

The authors wish to thank Professor A. Satyanarayana for many inspiring discussions with him.

Appendix

***Theorem* 1**: Let G be a layered–graph constructed from an optical network topology N and let M be a set of lightpath requests. Then, N can support all the lightpath requests in M as MCP if and only if G contains a set of edge–disjoint paths for all the requests in M.

Proof: Suppose N contains a set of MCPs for the lightpath requests in M. Then, each MCP of color i can be drawn with Wavelength i on the i^{th} layer of G, $i = 1, 2, .. |W|$. Since all the layers in G are disjoint, and since the MCPs which are assigned the same wavelength must be edge–disjoint in N, the paths drawn on the layered–graph are disjoint as well. Conversely, suppose that for the lightpath requests in M there exists a set of edge–disjoint paths in G. Then, since the layers in G are disjoint, Wavelength i can be assigned to all the paths on the i^{th} layer of G, $i = 1, 2, .. |W|$. Now, by merging nodes $v_i^1, v_i^2, .. v_i^W$ to a single node for each $i \in R$, all the layers of G are merged to a single layer. The source and destination nodes v_a^s, v_a^d are also merged to a single node for each $a \in A$. This one–layer graph is identical to network topology N where all the paths from all the layers will form a valid set of mono–chromatic paths satisfying the requests in M. □

References

[1] *IEEE/OSA Journal of Lightwave Technology,* special issue on Broad–Band Optical Networks, vol. 11, no. 5/6, May/June 1993.

[2] B. Mukherjee, "WDM–Bases Local Lightwave Networks: Part I: Single–Hop Systems, Part II: Multihop Systems," *IEEE Network*, pp. 12–27, May 1992 and pp. 20–32, July 1992.

[3] Z. Zhang and A. Acampora, "A Heuristic Wavelength Assignment Algorithm for Multihop WDM Networks with Wavelength Routing and Wavelength Reuse," *Proc. IEEE Infocom '94*, pp. 534–543, Toronto, June 1994.

[4] R. Ramaswami and K.N. Sivarajan, "Optimal Routing and Wavelength Assignment in All–Optical Networks," *Proc., IEEE Infocom '94,* pp. 970–979, Toronto, June 1994.

[5] K.–C. Lee and V.O.K. Li, "A Circuit Rerouting Algorithm for All–Optical Wide–Area Networks," *Proc. IEEE Infocom '94*, pp. 954–961, Toronto, June 1994.

[6] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath Communications: An Approach to High Bandwidth Optical networks," *IEEE Transactions on Communications,* vol. 40, no. 7, pp. 1171–1182, July 1992.

[7] R.K. Ahuja, K. Mehlhorn, J.B. Orlin, and R.E. Tarjan, "Faster algorithms for the shortest path problem," *Journal of the ACM*, Vol. 37, pp. 213–223.

[8] S. Banerjee and C. Chen, "Design of Wavelength–Routed Optical Networks for Circuit Switched Traffic," Tech. Report CS–9511, *Stevens Institute of Technology*, 1995.

[9] C. Chen and S. Banerjee, "A New Model for Optimal Routing and Wavelength Assignment in Wavelength Division Multiplexed Optical Networks," Tech. Report CS–9507, *Stevens Institute of Technology*, 1995.

[10] P. Raghavan, and E. Upfal, "Efficient Routing in All–Optical Networks," *Proc. 26th ACM STOC*, pp. 384–395, 1994.

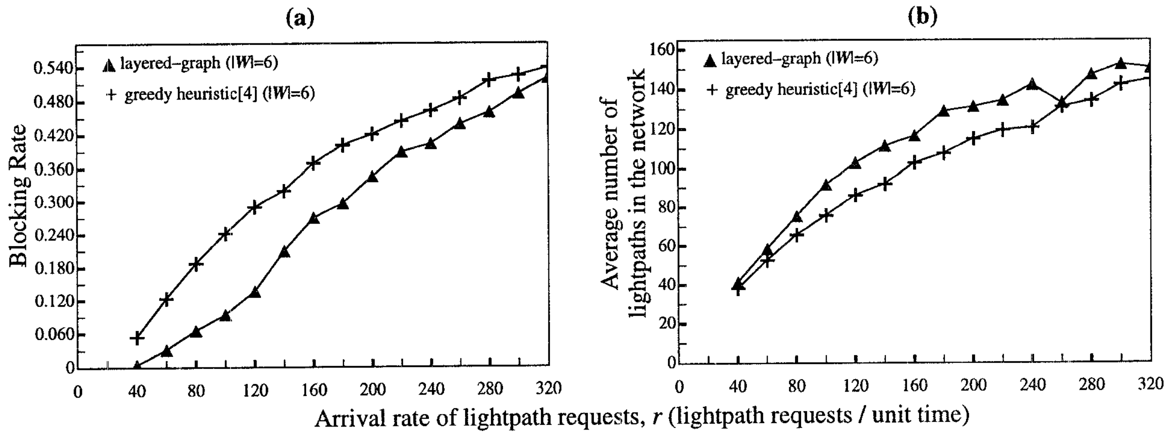

Fig. 5. Performance of the *dynamic RAW* algorithm: (a) blocking rate, (b) average number of lightpaths in the network, for different arrival rates of lightpath requests.

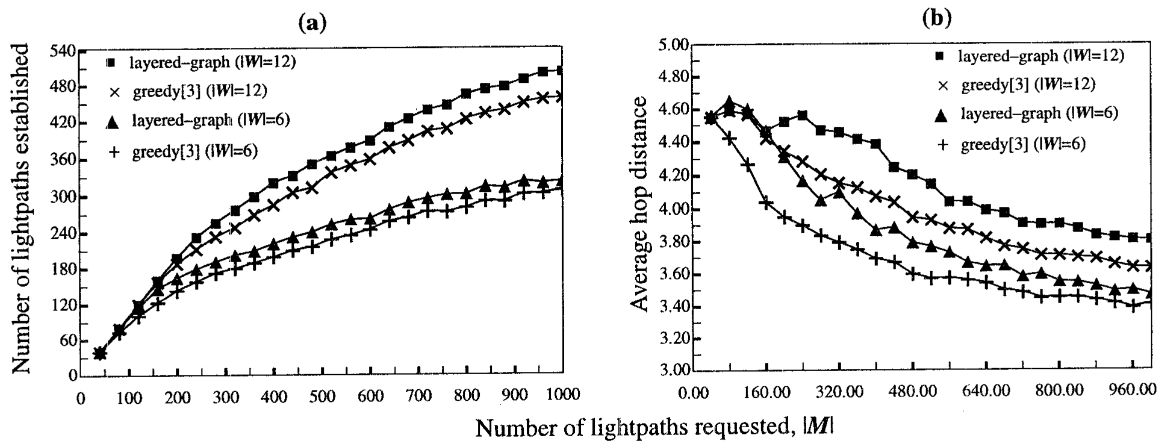

Fig. 6. Performance of *static RAW* algorithm for uniform traffic case: (a) number of lightpaths established, (b) average hop distance, for different number of lightpath requests.

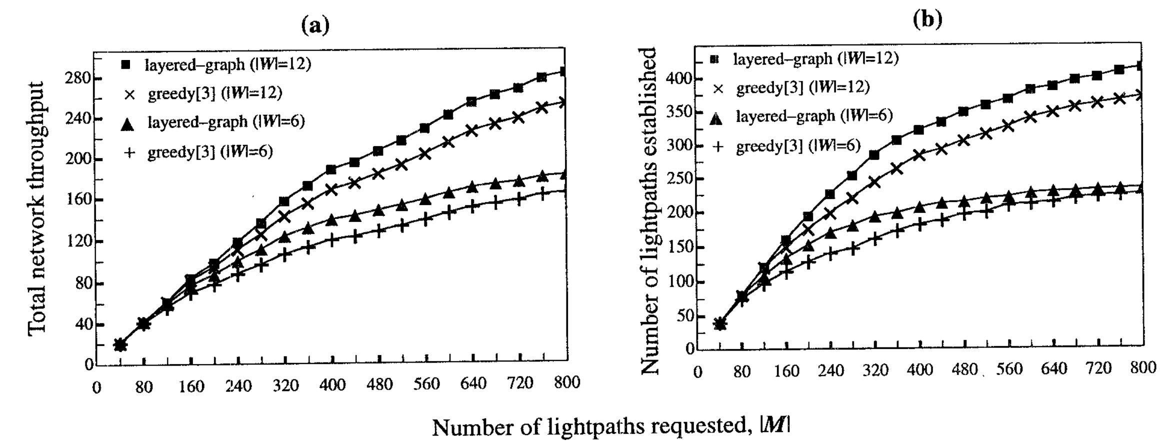

Fig. 7. Performance of *static RAW* algorithm for non–uniform traffic case: (a) total throughput of all the established lightpaths, (b) number of lightpaths established, for different number of lightpaths requested.