Computer Networks 52 (2008) 1891–1904

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Rerouting schemes for dynamic traffic grooming in optical WDM networks

Wang Yao [a,*], Byrav Ramamurthy [b]

[a] Blade Network Technologies Inc., 1901 Halford Ave. Apt 67, Santa Clara, CA 95051, United States
[b] Department of Computer Science and Engineering, University of Nebraska-Lincoln, Lincoln, NE 68588, United States

ARTICLE INFO

Article history:
Received 25 May 2007
Received in revised form 14 December 2007
Accepted 20 February 2008
Available online 29 March 2008

Keywords:
Rerouting
Traffic grooming
Dynamic traffic
Wavelength division multiplexing (WDM)

ABSTRACT

Traffic grooming in optical WDM mesh networks is a two-layer routing problem to effectively pack low-rate connections onto high-rate lightpaths, which, in turn, are established on wavelength links. The objective of traffic grooming is to improve resource efficiency. However, resource contention between lightpaths and connections may result in inefficient resource usage or even the blocking of some connections. In this work, we employ a rerouting approach to alleviate resource inefficiency and improve the network throughput under a dynamic traffic model. We propose two rerouting schemes, rerouting at lightpath level (RRLP) and rerouting at connection level (RRCON) and a qualitative comparison is made between the two. We also propose two heuristic rerouting algorithms, namely the critical-wavelength-avoiding one-lightpath-limited (CWA-1L) rerouting algorithm and the critical-lightpath-avoiding one-connection-limited (CLA-1C) rerouting algorithm, which are based on the two rerouting schemes. Simulation results show that rerouting reduces the blocking probability of connections significantly.

1. Introduction

Wavelength division multiplexing (WDM)-based optical networking architectures using optical cross-connects (OXCs) are promising solutions to the next-generation long-haul transport networks [1,2]. With an optical layer between the underlying point-to-point physical links and the upper electronic layer (SONET or IP/MPLS), this approach fully exploits the reconfiguration capability of today's optical equipment to provide flexible wavelength routing, protection, and restoration mechanisms. The transmission rate on a wavelength channel has reached OC-192 (10 Gbps) and is expected to reach OC-768 (40 Gbps) in the future. In dense WDM (DWDM) systems, a single fiber can carry more than 100 wavelengths. While WDM technology enables huge amounts of bandwidth, traffic demand is increasing at an explosive rate as well. The traffic demand granularity varies considerably, possibly from OC-3 (155 Mbps) to OC-192 (10 Gbps). The routing problem with the bandwidth gap between the low-rate connection and high-rate wavelength channels is addressed as a two-layer traffic grooming problem [3] with the goal of effectively sharing resources in the optical networks.

The two layers involved in traffic grooming are the optical layer and the electronic layer. The optical layer is composed of OXCs with point-to-point fiber links connecting them. The OXC nodes and the fiber links constitute the physical topology of an optical network. The optical layer establishes lightpaths along wavelength channels within the fibers. All the lightpaths and their corresponding end nodes constitute the virtual (logical) topology. The electronic layer establishes traffic connections on top of the virtual topology. Fig. 1 illustrates the two-layered routing in traffic grooming. As the optical layer and electronic layer functionally operate in different equipment, the simplest way to execute the routing of lightpaths in the optical layer and the routing of connections in the electronic layer is to execute them independently. However, studies [4,5] have indicated that considering the two-layer routing problem jointly achieves better resource efficiency.

* Corresponding author. Tel.: +1 408 313 0803.
E-mail addresses: wangyao@gmail.com (W. Yao), byrav@cse.unl.edu (B. Ramamurthy).

1389-1286/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2008.02.027

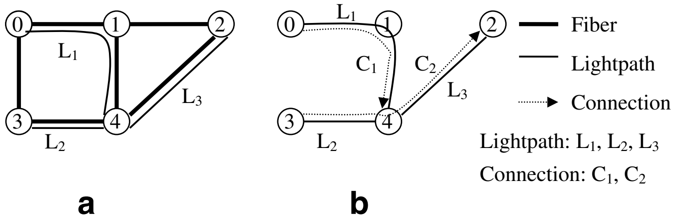

Fig. 1. Illustration of the two-layered routing in traffic grooming. (a) Physical topology of a five node network and three lightpaths established on the physical topology: L_1 on path 0–1–4, L_2 on path 3–4, L_3 on path 4–2. (b) Virtual topology and two connections established on the virtual topology: C_1 on lightpath L_1 and C_2 on path 3–4–2 using lightpaths L_2 and L_3.

Most of early work in traffic grooming aimed at minimizing electronic costs in SONET over WDM optical networks [5–9] while satisfying static traffic demands. Recent work in traffic grooming has focused more on the general WDM mesh optical network. The works in [10,11] formulated the static traffic grooming problem using integer linear programming (ILP) methods and presented heuristics to maximize network throughput given the network resources and static traffic demands. The work in [12] examined the survivable traffic grooming problem, which considered path protection at the connection level. ILP formulations and tabu-based heuristics were presented. The work in [13–15] studied the dynamic traffic grooming problem where traffic requests were assumed to arrive dynamically according to a random stochastic process. Two different graph-based grooming algorithms were proposed in [14,15]. The work in [16] studied multicast traffic grooming using a novel light-tree architecture. However, none of these works considered rerouting existing connections to satisfy new connections in dynamic traffic grooming.

Optimal routing in a two-layered optical network with dynamically arriving connection requests is a challenging task. Various factors may deteriorate the route selection of lightpaths and connections and thus reduce the total network throughput. For example, the resource contention between two lightpaths (or connections) may result in the blocking of one lightpath or the use of an inefficient long path for one of the two lightpaths. In a dynamic traffic environment, the contention is more a prominent problem, as the arrivals and departures of connections are not known in advance. Therefore, it is not possible to make the optimal routing decisions for all the lightpaths and connections at the same time so that the total network throughput can be maximized. As another example, consider optical networks without wavelength conversion capability. The wavelength continuity constraint stipulates that a lightpath must utilize the same wavelength along its path. This constraint reduces the possibility of successfully finding a free wavelength on a path and thus may force the lightpath to use another longer path or get blocked.

Rerouting is a useful technique to address the routing challenges mentioned above. While wavelength conversion is one technique to alleviate the inefficiency caused by wavelength continuity constraints, the work in [17] showed that rerouting also helps to increase the overall resource utilization efficiency. In [17], a move-to-vacant wavelength-retuning (MTV_WR) rerouting scheme was proposed. MTV means rerouting the lightpath to a path along which at least one wavelength is not occupied by any other lightpath. Using the MTV scheme, a new lightpath could be established on the new path before stopping the transmission on the old path and switching to the new path. Therefore, MTV achieves a small rerouting disruption time. On the other hand, WR means only retuning the wavelength of the lightpath on the same path, which makes the rerouting algorithm and operation simple. By combining the two schemes, MTV_WR tries to retune a lightpath to a free wavelength on the same path. It thus has the advantages of both MTV and WR. However, WR is not able to reduce the resource-inefficiency caused by lightpaths using extremely long paths. In [17], MTV_WR was used in the second phase of an overall two-phase routing scheme. The first phase was a normal routing process without considering rerouting. If the first phase failed to find a path for the lightpath request, then the MTV_WR rerouting was applied. In [18], a single-phase algorithm, still using the MTV_WR scheme, but combining routing and rerouting, was presented. It claimed to have less time complexity than the two-phase algorithm. Simulation results showed that it did run faster than the two-phase algorithm under moderate and heavy loads with high blocking probability. However, under relatively light loads with low blocking probability, which is more practical in reality, its speed was slower than the two-phase algorithm. This was because though the rerouting phase in the two-phase algorithm had a larger time complexity than the single-phase algorithm, it was rarely executed under relatively light load, where most requests have been satisfied in the first phase.

In this paper, we study the rerouting approach for the provisioning of multi-granularity connections in two-layer wavelength-routed optical networks with grooming capability. We consider the dynamic traffic environment where connection requests arrive and depart dynamically. As traffic grooming aims at improving the network resource utilization efficiency, we expect that using rerouting in traffic grooming will help achieve the objective. As in [17], we use a two-phase routing scheme. The rerouting procedure is applied only when the normal routing fails. We propose two rerouting schemes at two different levels, namely rerouting at the lightpath level (RRLP) and rerouting at the connection level (RRCON). To fully exploit the ability of rerouting to reduce the long path inefficiency, we use the idea of MTV, instead of MTV_WR, in the RRLP and RRCON rerouting schemes. We show that both schemes still preserve the short disruption time property while reducing the connection blocking probability. Two efficient heuristic rerouting algorithms based on RRLP and RRCON are proposed as well. To our knowledge, [21] is the only other work that has considered rerouting of existing connections in dynamic traffic grooming environments. The work in [21] focuses more on the rerouting of backup connections when path protection is provided. The other works in traffic grooming that have considered reconfiguration usually refer to adjusting the lightpaths in the virtual topology to accommodate more connections, and most of them assumed static traffic demands.

The remainder of the paper is organized as follows: Section 2 presents the two rerouting schemes, namely RRLP and RRCON. Section 3 presents two rerouting algorithms based on the two rerouting schemes. Section 4 compares the performance of the two rerouting algorithms by simulation in terms of connection blocking probability, rerouting frequency, rerouting success rate, percentage of traffic affected, and running time. Section 5 concludes the paper.

2. Rerouting schemes

A complete rerouting scheme generally has two components: a rerouting algorithm and a rerouting operation [17,18]. A rerouting algorithm determines whether in order to accommodate a new connection, existing lightpaths or connections will be rerouted, and if they are rerouted, which new paths they will use. The rerouting operation decides the sequence of steps executed in the networks to migrate the rerouted lightpaths or connections to their new paths. The rerouting operation belongs to the function of the control plane and largely determines the rerouting disruption time. For high-speed optical networks, even a short disruption time may affect a large amount of traffic. Therefore, it is desirable that a rerouting scheme should incur the minimum disruption time.

As the rerouting operation belongs to the functionality of the control plane, it can be implemented using generalized multi-protocol label switching (GMPLS) [19], which is expected to be the unified control plane for next-generation optical networks. GMPLS supports multiple switching types such as time-division multiplexing (TDM), wavelength (lambda), and fiber (port) switching. Accordingly, GMPLS supports a hierarchy of label switched paths (LSPs), which correspond to channels such as lightpaths at the wavelength level and SONET circuits at the subwavelength level.

In this section, we propose two rerouting schemes, RRLP and RRCON, to provide multi-granularity connections in grooming networks.

2.1. Rerouting at lightpath (RRLP) level

2.1.1. Rerouting algorithm

The basic idea of RRLP is to reroute some of the existing lightpaths so that new lightpaths can be established to carry an otherwise blocked connection. For example, for a network in Fig. 2a, assume that each fiber has only one wavelength λ_0 and a wavelength can carry two connections. Fig. 2b shows the current virtual topology and connections carried over the virtual topology. If a connection request 0–1 arrives, it cannot be accepted because neither an existing lightpath between node 0 and 1 exists nor a new lightpath 0–1 can be established. The wavelength λ_0 has been allocated to the lightpath 0–4, which uses path 0–1–4. However, if the lightpath 0–4 is rerouted to path 0–3–4, as shown in Fig. 2d, then a new lightpath 0–1 can be established to carry the new connection 0–1.

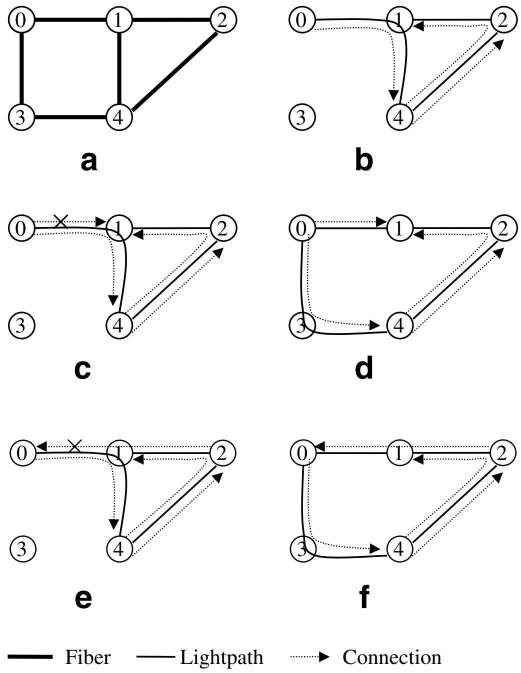

Fig. 2. Illustration of rerouting at lightpath level. Assume that each fiber has one wavelength λ_0 and each wavelength can accommodate two connections: (a) physical topology; (b) virtual topology and connections; (c) connection request 0–1 denied before rerouting; (d) connection request 0–1 accepted after rerouting; (e) connection request 2–0 denied before rerouting; (f) connection request 2–0 accepted after rerouting.

While the new lightpath established after rerouting can be the only lightpath needed to carry the connection, it is also possible that the new lightpath is one hop of the path to carry the connection. For example, in Fig. 2e, the connection request 2–0 is blocked because there is no lightpath between node 1 and 0 and the bandwidth capacity on lightpath 2–4 is used by two other connections. However, if the lightpath 0–4 can be rerouted to path 0–3–4, then a new lightpath 1–0 can be established, and the connection 2–0 can be carried over the path 2–1–0 using two lightpaths 2–1 and 1–0.

2.1.2. Rerouting operation

The lightpath rerouting operation is executed when the rerouting algorithm has decided which lightpaths are to be rerouted and what their new paths are. As multiple lightpaths may be rerouted, we assume that they are computed by the rerouting algorithm in such a way that they can be rerouted in parallel. This parallelism saves the time needed to wait for the completion of the rerouting process. The following procedure describes the basic steps to execute a lightpath rerouting operation. Once the rerouting operation finishes, the routing protocol can begin to establish the new connection.

For each lightpath to be rerouted:

1. Establish a new LSP for the rerouted lightpath.
2. Stop the data transmission on all connections using the rerouted lightpath.
3. Switch the rerouted lightpath from its old LSP to the new LSP.
4. Start the data transmission on all connections involved.
5. Release the wavelengths and transceivers used by the rerouted lightpath on the old LSP.

2.2. Rerouting at connection (RRCON) level

2.2.1. Rerouting algorithm

The basic idea of RRCON is to reroute some existing connections from certain lightpaths so that the otherwise blocked connection can use these lightpaths in its path. For example, for the network in Fig. 3a, assume that Fig. 3b shows its current virtual topology and connections. If a connection request 3–2 arrives, it will be blocked, because all the bandwidth capacity on lightpath 3–4 has been used by two other connections. However, if connection 0–4 can be rerouted to lightpath 0–4, which uses path 0–1–4, as shown in Fig. 3d, then connection 3–2 can be carried over a two-hop (virtual) path 3–4–2 using lightpaths 3–4 and 4–2.

2.2.2. Rerouting operation

The connection rerouting operation is executed when the rerouting algorithm has determined which connections are to be rerouted and their new paths. Note that in order to establish the new path for a rerouted connection, new lightpaths may need to be established.

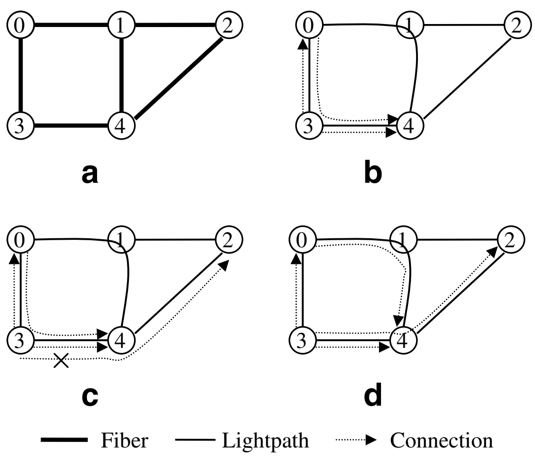

Fig. 3. Illustration of rerouting at connection level. Assume that each fiber has one wavelength λ_0 and each wavelength can accommodate two connections: (a) physical topology; (b) virtual topology and connections; (c) connection request 3–2 denied before rerouting; (d) connection request 3–2 accepted after rerouting.

For each connection to be rerouted:

1. Establish a new LSP for the rerouted connection. Establish new lightpaths if necessary.
2. Stop the data transmission on the rerouted connection.
3. Switch the rerouted connection from its old LSP to the new LSP.
4. Start the data transmission on the rerouted connection.
5. Release the lightpath bandwidth allocated in the old LSP of the rerouted connection.

2.3. Comparison of RRLP and RRCON

Generally, RRLP is a coarse-granularity rerouting scheme that operates at the high-rate lightpath (aggregate) level, while RRCON is a fine-granularity rerouting scheme that operates at the low-rate connection (per-flow) level. We compare RRLP and RRCON with respect to rerouting algorithm and rerouting operation respectively.

As a coarse-granularity scheme, the RRLP routing algorithm is relatively simple. It only needs the global information of the lightpaths in the network to decide the rerouted lightpaths and their new paths. The route changes of the underlying lightpaths can be viewed as being transparent to the upper connections. The RRCON routing algorithm, on the other hand, needs the global information of all the lightpaths and connections to compute the rerouted connections and their new paths. As the grooming factor C (which is equal to the capacity of a lightpath divided by the smallest allowed connection capacity) is usually a large number, the number of connections in the network is larger than the number of lightpaths. Therefore, the RRCON routing algorithm usually has a larger complexity. However, as a fine-granularity scheme, RRCON is more flexible in terms of selecting rerouted connections and their new paths. Also, as an end-to-end rerouting scheme with respect to connections, the RRCON routing algorithm may be good at preserving other quality of service (QoS) and traffic engineering (TE) constraints imposed on the connections.

From an operational point of view, RRLP usually affects more traffic than RRCON during the rerouting process. This is because in the LSP switch step, RRLP stops the data transmission on all the connections that use the rerouted lightpath, while RRCON only stops the data transmission on the rerouted connection. The service disruption time is low for both RRLP and RRCON. Theoretically, the service disruption time is equal to the time needed to switch from the old LSP to the new LSP. It is very short since the new LSP has been established before disrupting the data transmission.

3. Heuristic rerouting algorithms

Our proposed routing scheme for traffic grooming includes two phases: a normal routing phase and a rerouting phase. Note that in this paper, routing refers to a general concept that includes the grooming function. The rerouting phase is only executed when the normal routing phase fails to find a path for the arriving connection request. The rerouting phase tries to reroute some lightpaths or connections so that the connection request can be accepted. At the end of the rerouting phase, if some lightpaths or connections are rerouted successfully, then the connection request is provisioned again using the routing algorithm of the normal routing phase. We represent a connection request as $\Phi(s, d, t, \Delta_t, B)$, where s is the source node, d is the destination node, t is the arrival time of the traffic request, Δ_t is the required service time for this request, and B is the traffic bandwidth requirement.

In this section, we first propose the least virtual hop first (LVHF) routing algorithm for the normal routing phase. Then we propose two heuristic rerouting algorithms, critical-wavelength-avoiding one-lightpath-limited (CWA-1L) heuristic and critical-lightpath-avoiding one-connection-limited (CLA-1C) heuristic, which are used in the rerouting phase. The two heuristics are based on RRLP and RRCON respectively. As indicated by their names, CWA-1L and CLA-1C restrict the rerouting operation to one lightpath or connection each time for a connection request. This restriction not only reduces the complexity of the rerouting algorithms, but also reduces the amount of traffic affected by the rerouting operation.

3.1. Least virtual hop first (LVHF) routing algorithm

LVHF is based on the fixed alternate routing (FAR) approach [20]. LVHF tries to use the least number of virtual hops to carry a connection on the fixed alternate paths. Basically, LVHF has two steps. The first step is to select a physical path in the physical topology as a candidate path for the arriving connection request. The second step is to use existing lightpaths or establish new lightpaths on the candidate path to carry the connection. Depending on whether new lightpaths need to be established and the number of lightpaths used, the connection path can be divided into four categories, as illustrated in Fig. 4. Note that a (virtual) hop in a connection path is a lightpath in the virtual topology. To differentiate the connection path from the physical path in the physical topology, the term "route" is specially designated to represent a connection path in this paper.

- *Single-hop route using a new lightpath (SN-route).* For example, for a connection request from 0 to 4, suppose the route in Fig. 4b is the current candidate route. It is a SN-route if a new lightpath is established on the path 0–1–4 to satisfy the connection.
- *Single-hop route using an existing lightpath (SE-route).* For example, for a connection request from 0 to 2, suppose the route in Fig. 4c is the current candidate route. It is a SE-route if the existing lightpath L_1 on path 0–3–2 is used to satisfy the connection.
- *Multi-hop route using only existing lightpaths (MOE-route).* For example, for a connection request from 0 to 4, suppose the route in Fig. 4d is the current candidate route. It is a 2-hop MOE-route if the existing lightpaths L_1 and L_2 are used to satisfy the connection.
- *Multi-hop route using both new and existing lightpaths (MNE-route).* For example, for a connection request from 2 to 6, suppose the route in Fig. 4e is the current candidate route. It is a 2-hop MNE-route if a new lightpath on path 4–6 and the existing lightpath L_2 are used to satisfy the connection.

As shown in Fig. 4, there are different ways (candidate routes) to satisfy a connection on a physical path. For a physical path with H physical hops, there is one (C_0^{H-1}) single-hop route, C_1^{H-1} 2-hop routes, C_2^{H-1} 3-hop routes, ..., C_{H-1}^{H-1} H-hop route. The total number of candidate routes is $C_0^{H-1} + C_1^{H-1} + \cdots + C_{H-1}^{H-1} = 2^{H-1}$. As the number of candidate routes increases exponentially with H, it is impractical to search all the candidate routes in a large network. One way to reduce the number of candidate routes is to set up a limit on the number of virtual hops a route can take (H_v). For example, if we let $H_v = 2$, then an H-hop path have one (C_0^{H-1}) single-hop route and C_1^{H-1} 2-hop routes. The total number of candidate routes is $C_0^{H-1} + C_1^{H-1} = H$.

Suppose that each fiber has W wavelengths. T is the maximum number of lightpaths established on a path in the network, H_p is the maximum number of physical hops within a path in the network, H_v is the maximum number of virtual hops allowed within a route. Obviously, $T \leqslant W$ and $H_v \leqslant H_p$. The maximum number of candidate routes on a path is $R = C_0^{H_p-1} + C_1^{H_p-1} + \cdots + C_{H_v-1}^{H_p-1}$. If $H_v = 1$, $O(R) = O(1)$. If $H_v = 2$, $O(R) = O(H_p)$. If $H_v = 3$, $O(R) = O(H_p^2)$.

The LVHF algorithm is shown in Fig. 5. The k-shortest paths between all the node pairs are either fixed or calcu-

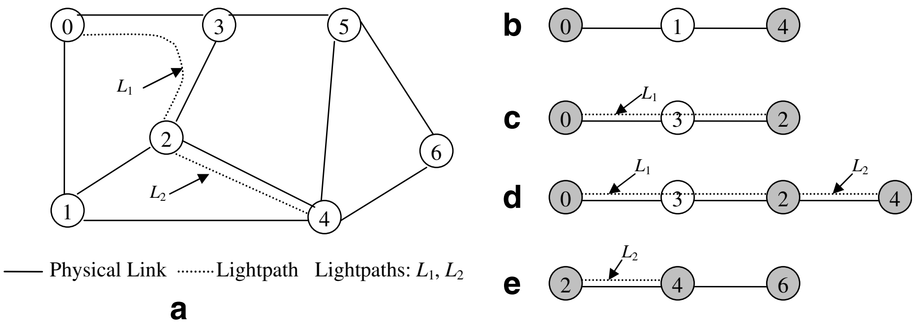

Fig. 4. Illustration of route categories: (a) the physical topology of a network and two lightpaths established on the physical topology; (b) a SN-route; (c) a SE-route; (d) an MOE-route; (e) an MNE-route.

```
Algorithm LVHF
Input: A connection request Φ and the current network state
information such as physical topology, virtual topology,
wavelength and transceiver usage and lightpath usage.
Output: A route in one of the four categories if accepted; NULL,
otherwise.
1) Set the k-shortest paths as the candidate paths P_c.
2) If an existing lightpath exists on one of the candidate paths
and has enough available bandwidth for Φ, return the lightpath as
a SE-route; otherwise, continue.
3) If a new lightpath can be established on one of the candidate
paths, return the new lightpath as a SN-route; otherwise, continue.
4) Search for a MOE-route or MNE-route for Φ.
     for h←2; h ≤ H_v ; h←h+1 do
          for each candidate path p ∈ P_c do
               for each h-hop route r on p do
4.1)                if ( Φ can be satisfied on the h-hop route r only
                        using existing lightpaths )
                       return the h-hop MOE-route r.
               for each h-hop route r on p do
4.2)                if ( Φ can be satisfied by the h-hop route r using
                       both existing lightpaths and new lightpaths )
                       return the h-hop MNE-route r.
5) return NULL.
```

Fig. 5. The LVHF algorithm.

lated periodically to reflect topology changes. Either way, the k-shortest paths should have been known before the provisioning of a specific connection request. Step 2 checks on each candidate path if there is an existing lightpath that can be used to carry the connection. Step 3 checks on each candidate path if new lightpaths can be established to carry the connection. For a wavelength, it needs to verify on all the links within the path if it is free. In the worst case, it needs to check all the W wavelengths. Step 4 tries to use multi-hop routes to carry the connection. Step 4.1 checks, on each virtual hop of a candidate route, whether an existing lightpath can be used to carry the connection. Step 4.2 checks whether an existing lightpath or a new lightpath can be used to carry the connection on each virtual hop of a candidate route.

Step 1 only takes O(1) time. Step 2 takes $O(kT)$ time, and step 3 takes $O(kWH_p)$ time. For step 4, the total number of candidate routes is $O(kR)$. Step 4.1 takes $O(H_vT)$ time. Step 4.2 takes $O(H_vT + H_pW)$ time. Therefore, step 4 takes $O(kR \times (2H_vT + H_pW))$ time. Because $T \leqslant W$ and $H_v \leqslant H_p$, it can be simplified to $O(kRH_pW)$. Hence the worst case time complexity of the LVHF algorithm is $O(1) + O(kT) + O(kWH_p) + O(kRH_pW) = O(kRH_pW)$.

3.2. CWA-1L rerouting heuristic

Definition 1. The critical wavelength set CW(p) of a path p is determined as given in (1). A wavelength is in CW(p) if it is used by a lightpath on only one fiber link of the path p and it is free on the rest of the links. The lightpath using the critical wavelength λ is called a critical wavelength related lightpath (CWR-lightpath).

$$\mathrm{CW}(p) = \{\lambda : \lambda \in \Lambda \text{ and } \exists e \in p \exists l \in \mathrm{LP}\ \lambda(e) \in l \text{ and } \forall e' \in p - e \forall l \in \mathrm{LP}\ \lambda(e') \notin l\}, \tag{1}$$

where Λ is the wavelength set supported on each fiber, $\lambda(e)$ is the wavelength λ on the fiber link e, e' is a fiber link on path p excluding e, LP is the set of all the lightpaths in the network, and $\lambda \in l$ means that lightpath l uses the wavelength λ.

According to the definition of CW(p), once a wavelength $\lambda \in \mathrm{CW}(p)$ becomes free, a new lightpath can be established on p using the wavelength λ. Since any $\lambda \in \mathrm{CW}(p)$ is used by a lightpath on a certain link, we have to use the rerouting technique to make λ free. If we can reroute a CWR-lightpath of path p without using its corresponding critical wavelength λ, then a new lightpath can be established on path p to satisfy the new connection. Fig. 6 shows the CWA-1L heuristic algorithm and the NEW_LP_CWA procedure used in CWA-1L.

The NEW_LP_CWA procedure basically tries to reroute a lightpath to make a critical wavelength become free. In step 1, NEW_LP_CWA identifies the critical wavelength set by checking hop by hop whether a wavelength is used by a lightpath. In step 2, NEW_LP_CWA tries to reroute a CWR-lightpath without using the corresponding critical wavelength.

In CWA-1L, step 2 tries to satisfy a connection with a SN-route by rerouting a CWR-lightpath corresponding to a critical wavelength of a candidate path. Step 3 tries to satisfy a connection with a MNE-route by rerouting a CWR-lightpath corresponding to a critical wavelength associated with a sub-path of a candidate path. For each candidate route, step 3.1 checks each virtual hop in the route whether an existing lightpath or a new lightpath can be used to satisfy the connection. CWA-1L only allows

```
Algorithm CWA-1L
Input: A connection request Φ and the current network state
information such as physical topology, virtual topology,
wavelength and transceiver usage and lightpath usage.
Output: A lightpath to be rerouted and the new path of the
rerouted lightpath, if successfully; NULL, otherwise.
1) Set the k-shortest paths as the candidate paths P_c.
2) Try to satisfy Φ with a SN-route using rerouting:
     for each candidate path p ∈ P_c do
          if ( NEW_LP_CWA(p, l, p*) )
               return l and p*.
3) Try to satisfy Φ with a MNE-route using rerouting:
     for h←2; h ≤ H_v ; h←h+1 do
          for each candidate path p ∈ P_c do
               for each h-hop candidate-route r obtained by dividing
               p into h sub-paths do
3.1)                if ( there is only one sub-path p_s on which neither
                       an existing lightpath with enough bandwidth
                       exists nor a new lightpath can be established )
3.2)                     if ( NEW_LP_CWA(p_s, l, p*) )
                              return l and p*.
4) return NULL.
```

```
Procedure NEW_LP_CWA(p, l, p*)
Input: A path p on which a new lightpath needs to be established.
Output: A lightpath l to be rerouted and the new path p* of l.
Return NULL if no lightpath can be rerouted.
1) Identify the critical wavelength set CW(p).
2) for each critical wavelength λ ∈ CW(p) do
        Identify the CWR-lightpath l of λ.
        if ( a new path p* without using λ is found for l )
             return l and p*.
3) return NULL.
```

Fig. 6. The CWA-1L algorithm and the NEW_LP_CWA method.

the rerouting of one lightpath, and only if a route has one hop (sub-path) on which neither an existing lightpath nor a new lightpath can carry the connection, does CWA-1L call the NEW_LP_CWA procedure to reroute a CWR-lightpath associated with the sub-path.

For the NEW_LP_CWA procedure, step 1 takes $O(H_pW)$ time. For step 2, the maximum number of wavelengths in a critical wavelength set is W. Identifying the CWR-lightpath of a critical wavelength only takes $O(1)$ time and finding a new path and a wavelength for a lightpath takes $O(kW)$ time. Therefore, step 2 takes $O(W(1+kW))$ time. The worst case time complexity of the NEW_LP_CWA procedure is $O(H_pW)+O(W(1+kW))=O(W(H_p+kW))$. In a network with tens of wavelengths, we would expect that $H_p < kW$. Therefore, the complexity can be simplified to $O(kW^2)$.

For the CWA-1L algorithm, step 2 takes $O(k^2W^2)$ time. As with LVHF, step 3 has $O(kR)$ candidate routes. Step 3.1 takes $O(H_vT+H_pW)=O(H_pW)$ time. Therefore, the worst case time complexity of step 3 is $O(kR)\times(O(H_pW)+O(kW^2))=O(k^2RW^2)$. The worst case time complexity of CWA-1L is $O(k^2W^2)+O(k^2RW^2)=O(k^2RW^2)$.

3.3. CLA-1C rerouting heuristic

Definition 2. The critical lightpath set $CL(p,\Phi)$ of a path p and a connection request Φ is determined as given in (2). A lightpath l on the path p is in $CL(p,\Phi)$ if it will have enough bandwidth to satisfy the connection request Φ only after releasing the bandwidth allocated to another connection c. Correspondingly, c is a critical lightpath related connection (CLR-connection) of the critical lightpath l. Note that a critical lightpath may have multiple CLR-connections. For a critical lightpath l associated with path p and request Φ, all its CLR-connections constitute the CLR-connection set $CLRC(p,\Phi,l)$.

$$CL(p,\Phi)=\{l\in LP: l \text{ is on path } p\}\cap\{l\in LP: B_a(l)<B_\Phi \text{ and } \exists c\in CON(l) B_c+B_a(l)\geqslant B_\Phi\}, \tag{2}$$

where $B_a(l)$ is the available bandwidth on the lightpath l, B_Φ is the bandwidth requirement of the connection request Φ, $CON(l)$ is the set of connections that use lightpath l, and B_c is the bandwidth of the connection c.

According to the definition of $CLRC(p,\Phi,l)$, if any CLR-connection is rerouted without using the corresponding critical lightpath l, then l will have enough free bandwidth capacity to carry the new connection request Φ. Fig. 7shows the CLA-1C heuristic algorithm and the FREE_CON_CLA procedure used in CLA-1C.

The FREE_CON_CLA procedure first identifies critical lightpaths on the path p and their corresponding CLR-connections. It then tries to reroute one CLR-connection without using its corresponding critical lightpath. Step 1 identifies critical lightpaths by checking each lightpath on path p to determine whether any connection c using the lightpath can satisfy the condition defined in (2). For each critical lightpath, step 2.1 checks all connections using this lightpath to identify its CLR-connections. Then step 2.3 tries to find a new route for an identified CLR-connection without using the critical lightpath. This can be accomplished by using any grooming algorithm, and marking the critical lightpath as temporarily unavailable. In this paper, step 2.3 of FREE_CON_CLA uses the modified LVHF algorithm, which only allows the rerouted connection to use existing lightpaths.

Algorithm CLA-1C

Input: A connection request Φ and the current network state information such as physical topology, virtual topology, wavelength and transceiver usage, lightpath usage and the connections.
Output: A connection to be rerouted and the new route of the rerouted connection, if successfully; NULL, otherwise.
1) Set the k-shortest paths as the candidate paths P_c.
2) Try to satisfy Φ with a SE-route using rerouting:
  for each candidate path $p\in P_c$ **do**
    if (FREE_CON_CLA(p, Φ, c, r^*))
      return c and p^*.
3) Try to satisfy Φ with a multi-hop route using rerouting:
  for $h\leftarrow 2$; $h\le H_v$; $h\leftarrow h+1$ **do**
    for each candidate path $p\in P_c$ **do**
      for each h-hop candidate-route r obtained by dividing p into h sub-paths **do**
3.1)   **if** (only one sub-path p_s on which an lightpath exists but without enough free bandwidth for Φ, and for the rest of the sub-paths, either a new lightpath can be established or an existing lightpath has enough free bandwidth)
3.2)   **if** (FREE_CON_CLA(p_s, Φ, c, r^*))
            return c and p^*.
4) **return** NULL.

Procedure FREE_CON_CLA(p, Φ, c, r^*)

Input: A path p on which an existing lightpath needs to reroute one of its connections so that the lightpath has enough free bandwidth for Φ.
Output: A connection c to be rerouted and the new route r^* of c. Return NULL if no connection can be rerouted.
1) Identify the critical lightpath set $CL(p, \Phi)$.
2) **for** each critical lightpath $l\in CL(p,\Phi)$ **do**
2.1) Identify the CLR-connection set $CLRC(p, \Phi, l)$.
2.2) **for** each CLR-connection c in $CLRC(p, \Phi, l)$ **do**
2.3)   **if** (a new route r^* without using l is found for c)
      return c and r^*.
3) **return** NULL.

Fig. 7. The CLA-1C algorithm and the FREE_CON_CLA method.

In CLA-1C, step 2 tries to reroute a connection on an existing lightpath on the path p so that the arriving connection request can be satisfied on the existing lightpath. Step 3 tries to reroute a connection on an existing lightpath on a sub-path of the path p so that the arriving connection request can be satisfied on a multi-hop route that includes the existing lightpath as a hop. For each candidate route, step 3.1 checks each hop (sub-path) to determine whether a new lightpath or an existing lightpath can satisfy the connection. If there is only one hop that cannot satisfy the connection using either a new lightpath or an existing lightpath, and there are existing lightpaths on the hop, then FREE_CON_CLA is called to reroute a CLR-connection to make room for the arriving connection request.

To analyze the time complexity of CLA-1C, we denote the maximum number of critical lightpaths on a path as T, which is the maximum number of lightpaths on a path. Let C be the maximum number of connections a lightpath can accommodate. In FREE_CON_CLA, step 1 takes $O(TC)$

time, and step 2.1 takes O(C) time. Step 2.3 is the modified LVHF algorithm which does not have steps 3 and 4.2 in LVHF. Therefore, its time complexity is O(kRH_v T). The worst case time complexity of FREE_CON_CLA is $O(TC) + O(T) \times (O(C) + O(C) \times O(kRH_vT)) = O(kRH_vT^2C)$.

For the CLA-1C algorithm, step 2 takes $O(k) \times O(kRH_vT^2C) = O(k^2RH_vT^2C)$ time. Step 3.1 takes $O(H_vT + H_pW) = O(H_pW)$ time. As with LVHF, there are O(kR) candidate routes for step 3. Therefore, the worst case time complexity of step 3 is $O(kR) \times (O(H_pW) + O(kRH_vT^2C)) = O(kRH_pW + k^2R^2H_vT^2C)$. The worst case time complexity of CLA-1C is also $O(kRH_pW + k^2R^2H_vT^2C)$.

3.4. Time complexity comparison

Table 1 shows the worst case time complexity of the three algorithms. In a WDM optical network with tens of wavelengths, we would expect that $H_p < kW$, which leads to the conclusion that $O(kRH_pW) < O(k^2RW^2)$. Therefore, the time complexity of CWA-1L is higher than that of LVHF.

The time complexity of CLA-1C is at least as high as, if not higher than, that of LVHF because of the common term O(kRH_pW) included in their complexity. For multi-hop traffic grooming, $H_v \geqslant 2$. Therefore, $O(R) \geqslant O(H_p)$. If $O(W) < O(kH_vT^2C)$, which is very likely, we can get $O(kRH_pW) < O(k^2R^2H_vT^2C)$. In this case, the time complexity of CLA-1C is higher than that of LVHF.

To compare the complexity of CWA-1L and CLA-1C, we need to compare $O(k^2RW^2)$ and $O(k^2R^2H_vT^2C)$. As $O(k^2RW^2)/O(k^2R^2H_vT^2C) = O(W^2)/O(RH_vT^2C)$, it is hard to determine which algorithm has a higher time complexity. In fact, it depends on parameters such as the number of wavelengths (W), the number of virtual hops allowed in a route (H_v), and the number of connections a lightpath can accommodate (C). To make the comparison more complicated, the parameters are correlated. For example, the maximum number of lightpaths that can be established on a path (T) is related to W. T must not be larger than W and should be small relative to W. For another example, the candidate routes search space (R) depends on the number of virtual hops allowed in a route (H_v). Generally, the complexity of CWA-1L heavily depends on W and the complexity of CLA-1C depends on a unique parameter C. The real average running times of CWA-1L and CLA-1C will be compared in Section 4.

4. Numerical results

To evaluate the performance of the two proposed rerouting heuristics, we simulate them on the NSFNET and EUPAN networks (shown in Fig. 8). The following assumptions are used. Connection requests arrive in a Poisson process with rate λ and are uniformly distributed among all the node pairs. The connection service time is distributed exponentially with mean $1/\mu$. Suppose the capacity (C) of a lightpath is normalized by dividing it by the smallest grooming granularity. Without explicit specification, we assume C is 16 and the connection bandwidth requirement is uniformly distributed between 1 and $C/2$. Each fiber supports 16 wavelengths and each node in the network has 32 transceivers. In this section, LVHF denotes the LVHF routing algorithm without the rerouting phase. CWA-1L and CLA-1C denote using LVHF in the normal routing phase and using the CWA-1L and CLA-1C algorithms respectively in the rerouting phase.

Table 1
The worst case time complexity of the algorithms

Algorithms	LVHF	CWA-1L	CLA-1C
Complexity	$O(kRH_pW)$	$O(k^2RW^2)$	$O(kRH_pW + k^2R^2H_vT^2C)$

W is the number of wavelengths on each fiber, k is the number of alternate paths, T is the maximum number of lightpaths established on a path, H_p is the maximum number of physical hops within a path, H_v is the maximum number of virtual hops allowed within a route, R is the maximum number of candidate routes on a path, and C is the maximum number of connections a lightpath can accommodate.

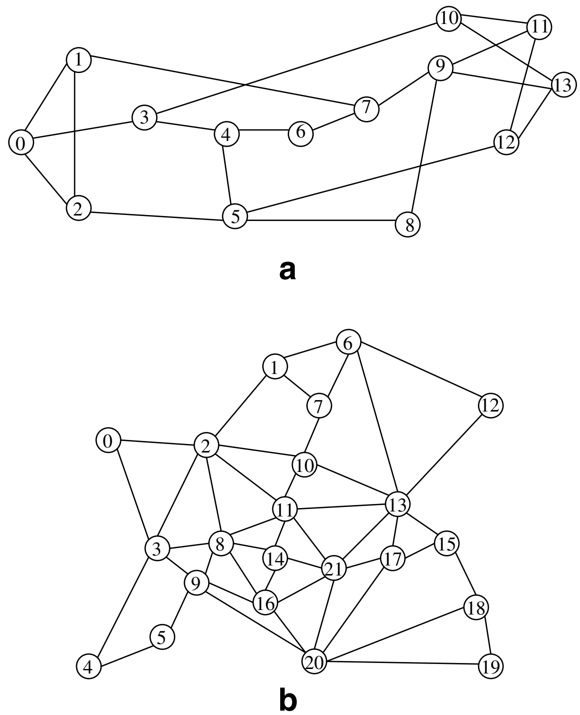

Fig. 8. Networks for simulation: (a) NSFNET network; (b) EUPAN network.

4.1. Blocking performance

Fig. 9 compares the performance of CWA-1L and CLA-1C with LVHF on the NSFNET network in terms of connection request blocking probability. We can see that with the same k (the number of alternate paths), CWA-IL and CLA-1C always yield a lower blocking probability than LVHF. The only exception is when $k = 1$, CWA-1L does not improve much over LVHF. This is because when we restrict the candidate path to be the shortest path, not many choices are left for CWA-1L to reroute a lightpath. The only possibility is to switch to a different wavelength on the same path. However, this effort seems to be futile according to the simulation results. CLA-1C, on the other hand, performs better than LVHF when $k = 1$. This is because that

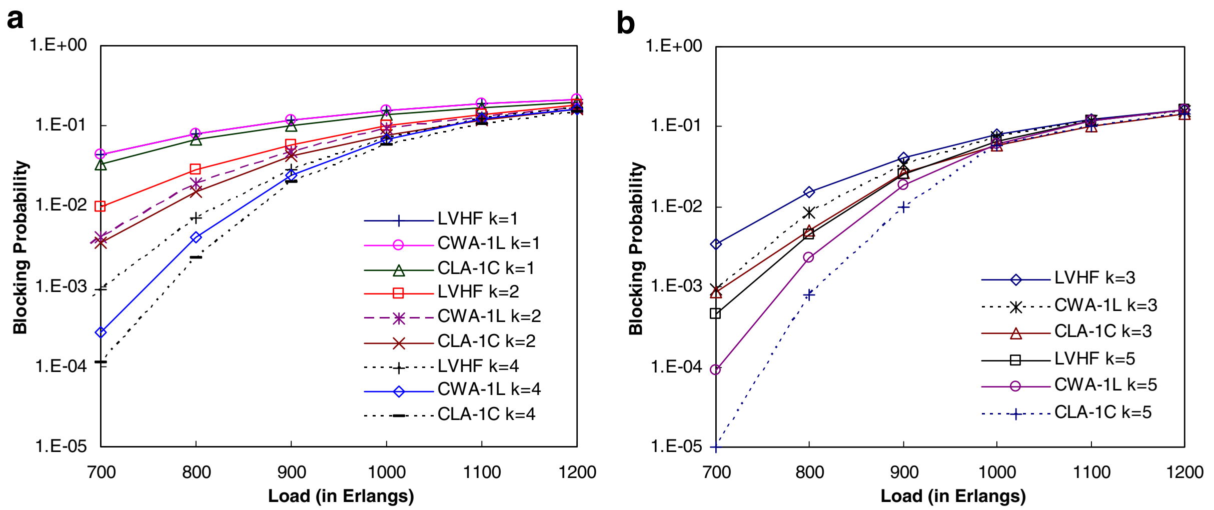

Fig. 9. Blocking performance of the rerouting algorithms in the NSF network: (a) k (the number of alternate paths) is 1, 2 and 4 respectively; (b) k is 3 and 5.

even though the physical path is restricted to the shortest path, CLA-1C can still use different kinds of lightpaths concatenations to reroute a connection. We can also see that CLA-1C performs better than CWA-1L on the NSFNET network when k is varied from 1 to 5.

Fig. 10 shows the simulation results on the EUPAN network. As can be seen, CWA-1L and CLA-1C perform better than LVHF when k is varied from 2 to 5. This result, along with the result from Fig. 9, substantiates our claim that rerouting helps to increase the network throughput, and thus enhances the resource utilization. Comparing CWA-1L and CLA-1C, we can see that CWA-1L yields lower blocking probability on the EUPAN network than CLA-1C when k is 2 or 3. When k becomes larger than 3, CWA-1L only performs better than CLA-1C under extremely light traffic load and is outperformed by CLA-1C under relatively medium or heavy traffic load. The reason why CWA-1L performs relatively worse than CLA-1C under heavy loads is that almost all wavelengths are consumed by other lightpaths in such cases. CWA-1L can hardly find available resources to reroute a lightpath. As for CLA-1C, it is more flexible in utilizing the available bandwidth in the lightpaths to reroute a relatively low-rate connection. Also, the reason why CWA-1L performs relatively worse at large k is that a large k may result in more inefficient long paths being used by lightpaths. Though it may result in accepting a few specific connections, it reduces the resource efficiency overall. This problem does not exist with CLA-1C because CLA-1C operates at much lower granularities than CWA-1L. A small proportion of the connections using relatively long paths do not significantly reduce the resource efficiency. On the contrary, it reduces the blocking probability by accepting the otherwise rejected connection requests.

By comparing Figs. 9 and 10, we can see that CWA-1L performs better in a relatively dense network like EUPAN, rather than in a relatively sparse network like NSFNET. This is because in a dense network there are many alternative paths with lengths close to or equal to the length of the shortest path. Therefore, rerouting a lightpath to an alternative path does not reduce the overall resource efficiency. However, in a sparse network, the alternative paths may be far longer than the shortest path. Rerouting a lightpath to a longer alternative path may reduce the overall resource efficiency, which in turn may compromise the effect of rerouting.

4.2. Rerouting frequency and success rate

This section compares the performance of CWA-1L and CLA-1C in terms of rerouting frequency (RRF) and rerouting success rate (RRSR). RRF, as shown in (3), is defined as the percentage of connection requests (REQs) that invoke the rerouting phase after not being satisfied in the normal routing phase:

$$\mathrm{RRF} = \frac{\text{number of REQs invoking rerouting phase}}{\text{total number of REQs}}. \tag{3}$$

RRSR, as shown in (4), is defined as the percentage of connection requests that are satisfied in the rerouting phase:

$$\mathrm{RRSR} = \frac{\text{number of REQs satisfied in rerouting phase}}{\text{number of REQs invoking rerouting phase}}. \tag{4}$$

Fig. 11 shows the rerouting frequency of CWA-1L and CLA-1C. As can be seen, CWA-1L and CLA-1C have similar rerouting frequencies, which should roughly be equal to the blocking probability of the LVHF algorithm. However, the rerouting phase will affect the normal routing phase. Even though CWA-1L and CLA-1C use the same LVHF algorithm in the normal routing phase, their rerouting frequencies are different. Generally, CLA-1C has a higher rerouting frequency than CWA-1L. This is probably because in the rerouting phase CLA-1C does not establish new lightpaths and CWA-1L always establishes new lightpaths to satisfy a connection request. Therefore, for the next-round routing phase, CLA-1C would have less available bandwidth left in the existing lightpaths than does CWA-1L.

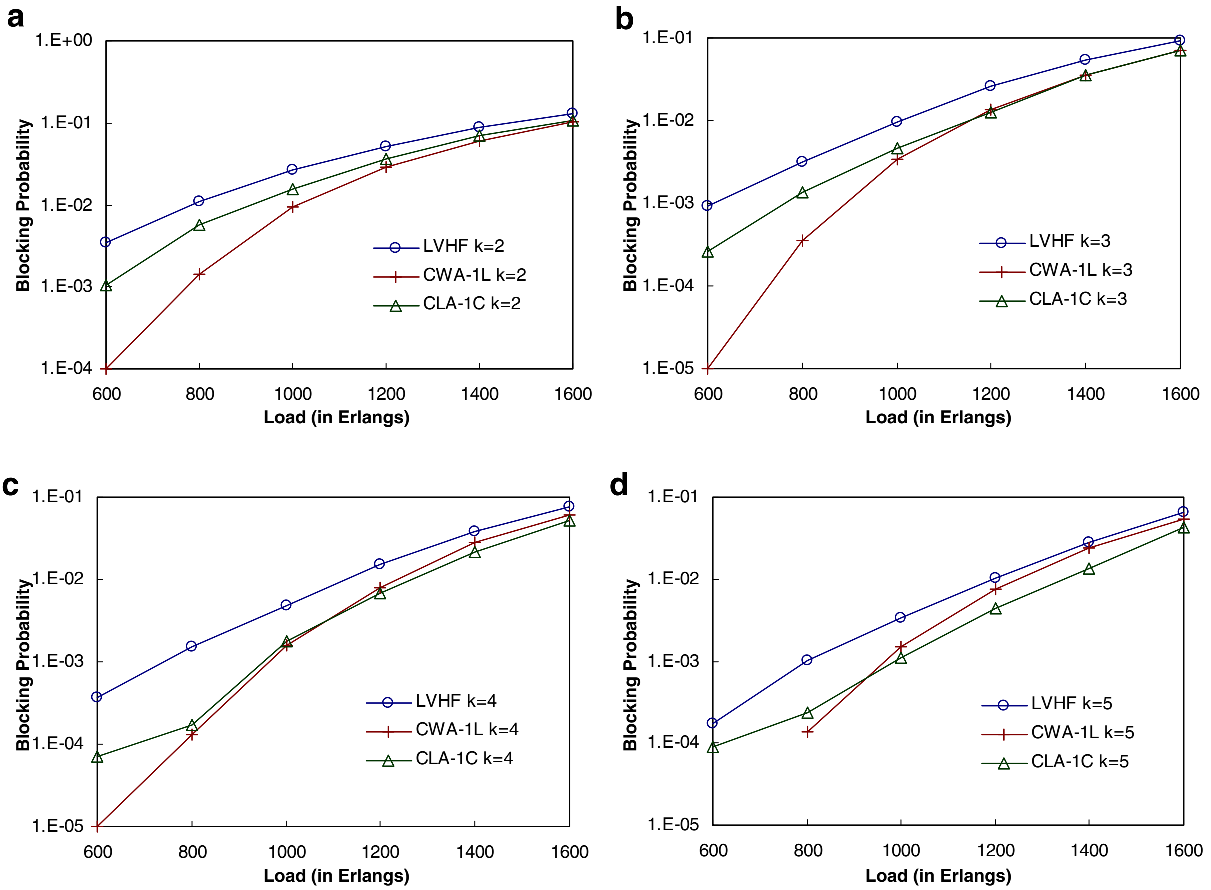

Fig. 10. Blocking performance of the rerouting algorithms in the EUPAN network: (a) k (the number of alternate paths) = 2; (b) k = 3; (c) k = 4; (d) k = 5.

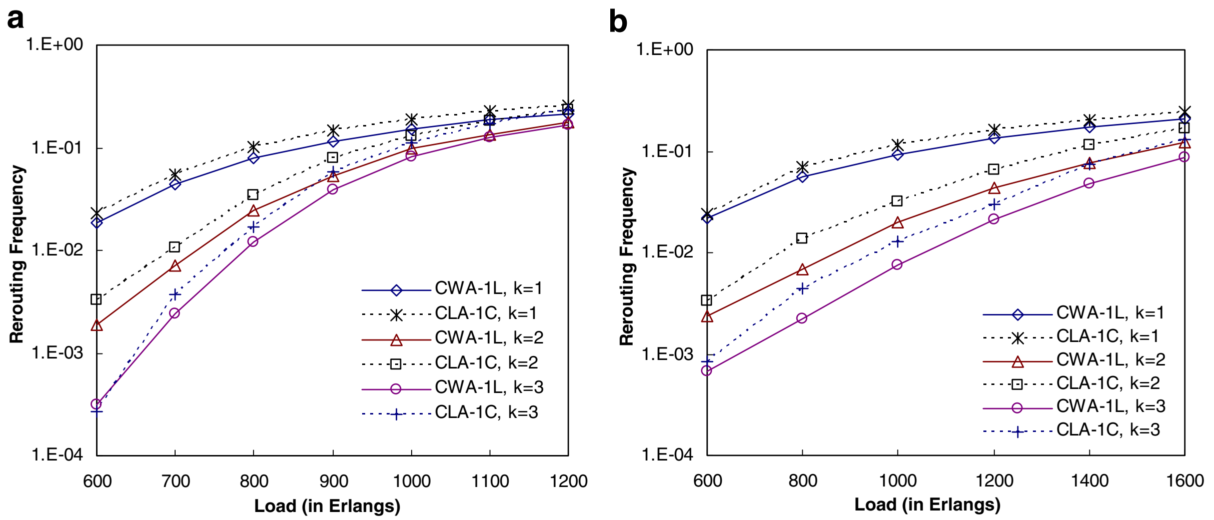

Fig. 11. Rerouting frequency of the rerouting algorithms: (a) NSF network; (b) EUPAN network.

Fig. 12 shows the rerouting success rates of CWA-1L and CLA-1C. As can be seen, the rerouting success rates of both CWA-1L and CLA-1C increase with k. This is reasonable because a larger k means more alternate paths, which would increase the chance of a successful rerouting. Specially, when $k = 1$, the rerouting success rate of CWA-1L

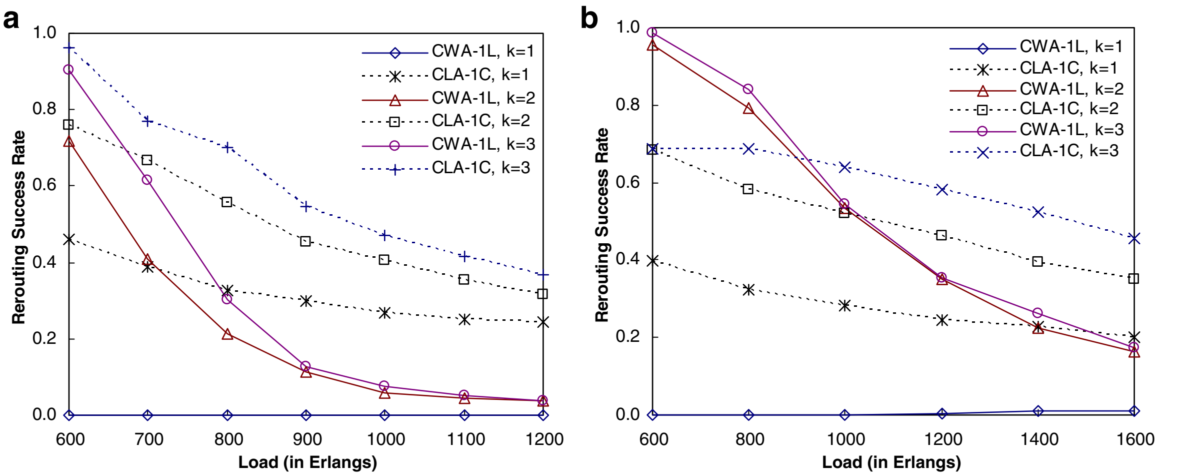

Fig. 12. Rerouting success rate of the rerouting algorithms: (a) NSF network; (b) EUPAN network.

is almost always 0 due to very limited rerouting choices. This explains why LVHF and CWA-1L have almost the same curve in Fig. 9a when $k = 1$.

A notable fact is that the rerouting success rate of CWA-1L drops quickly with the increase of the traffic load. This is because that under heavy loads, there would be few resources left in the network for CWA-1L to establish new lightpaths. Therefore, the chance of a successful rerouting decreases. Another interesting result is that the rerouting success rate of CWA-1L does not drop as quickly in the relatively dense EUPAN network as it does in the relatively sparse NSF network. On the other hand, CLA-1C is a more flexible rerouting algorithm as it is at a fine-granularity connection level. Therefore, its rerouting success rate does not drop as dramatically as CWA-1L with the increase of traffic load. This explains why CLA-1C can still have a better blocking performance than CWA-1L even though CLA-1C has a larger rerouting frequency than CWA-1L.

4.3. Rerouting affected traffic

According to the qualitative analysis in Section 2, the rerouting algorithms based on the RRLP and RRCON rerouting schemes should have short disruption times for the connections being rerouted. In this section, we inspect the total amount of traffic affected (rerouted) by CWA-1L and CLA-1C. Obviously, CWA-1L has a larger rerouting granularity than CLA-1C since a rerouting in CWA-1L affects all the connections within rerouted lightpaths, and a rerouting in CLA-1C only affects the rerouted connections. However, this does not necessarily lead to the conclusion that CWA-1L will affect more traffic than CLA-1C. Since CLA-1C generally has a higher rerouting frequency and higher rerouting success rate than CWA-1L, it is possible that the total amount of traffic affected by CLA-1C is larger than that of CWA-1L. As shown in Fig. 13a, CLA-1C affects more traffic than CWA-1L under relatively heavy

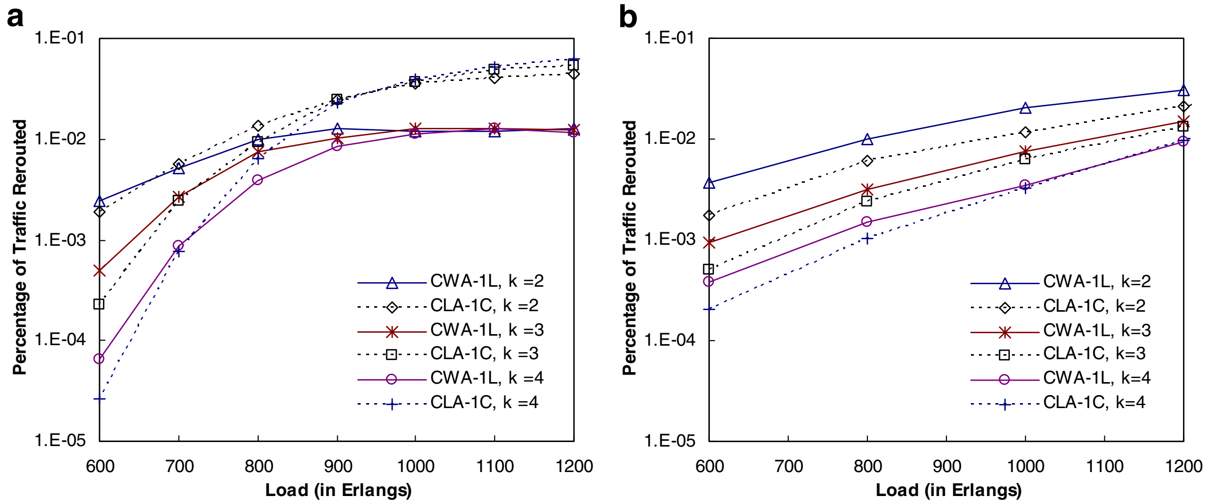

Fig. 13. Percentage of traffic affected: (a) NSF network; (b) EUPAN network.

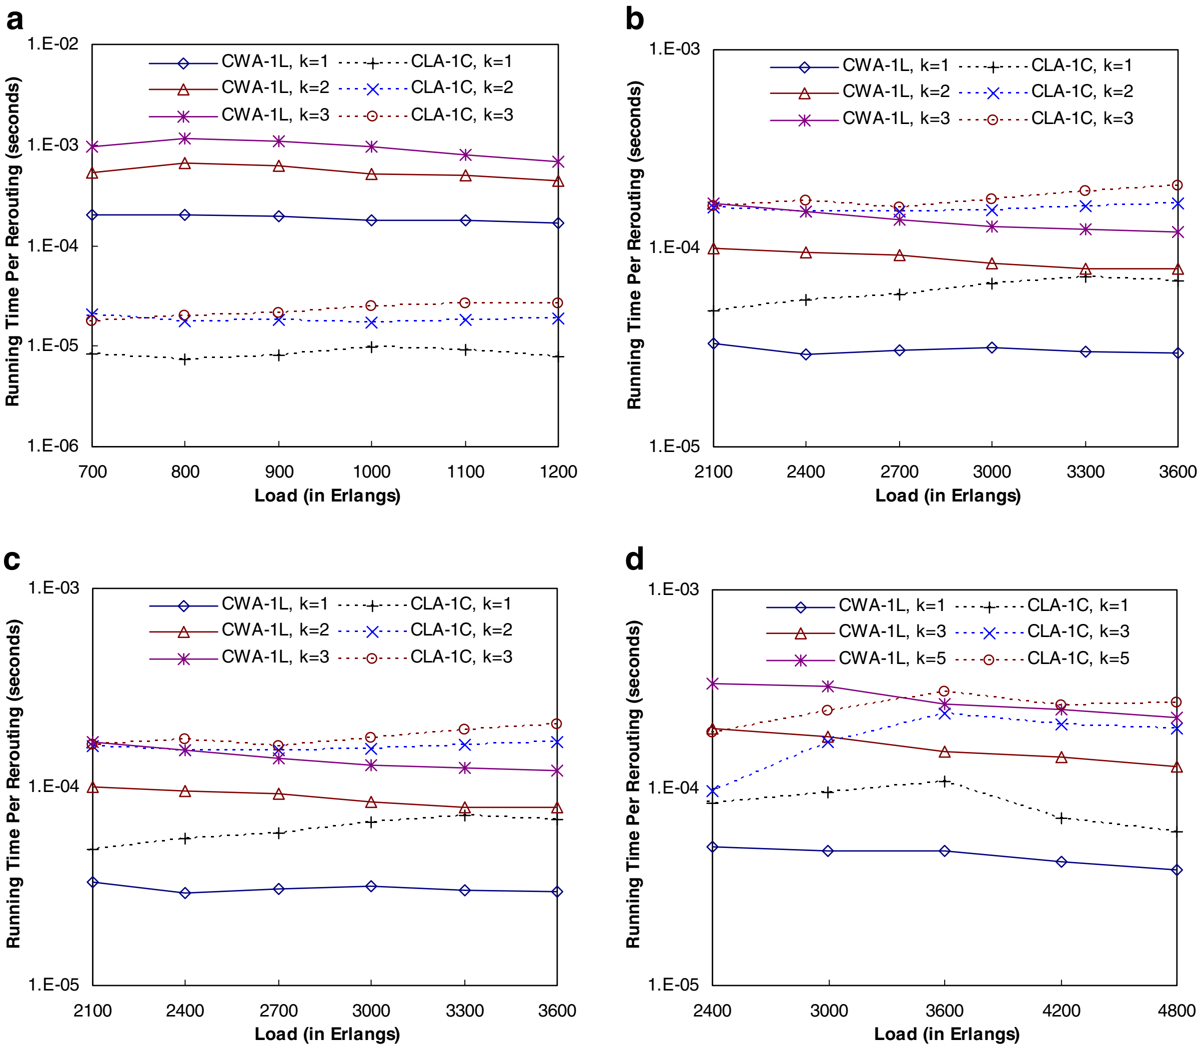

Fig. 14. Average running time per rerouting: (a) NSF network, $W = 16$, $C = 16$, $G = 8$; (b) NSF network, $W = 4$, $C = 64$, $G = 2$; (c) EUPAN network, $W = 16$, $C = 16$, $G = 8$; (d) EUPAN network, $W = 4$, $C = 64$, $G = 2$.

loads in NSF network. In Fig. 13b, CWA-1L affects more traffic than CLA-1L in the EUPAN network. This is because CWA-1L has a higher rerouting success rate in relatively dense networks and under relatively low traffic loads.

4.4. Running time

As discussed in Section 3.4, the worst-case time complexity of the two algorithms depends on some different parameters. Therefore, it is hard to determine which is more efficient. Fig. 14 compares the average running times per rerouting of CWA-1L and CLA-1C under different scenarios. Specifically, we consider two different scenarios. In scenario 1, the number of wavelengths (W) is relatively large and the number of connections a lightpath can accommodate is relatively small. Scenario 2 is the opposite of scenario 1. Figs. 14a and b are the results of the NSF network. Fig. 14a simulates scenario 1, where $W = 16$, the bandwidth capacity of a lightpath (C) is 16 units and the maximum bandwidth of a connection (G) is 8 units. Fig. 14b simulates scenario 2, where $W = 4$, $C = 64$, and $G = 2$. For scenario 1, CWA-1L is slower than CLA-1C. For scenario 2, CWA-1L is faster than CLA-1C. Similar results are shown in Figs. 14c and d for the EUPAN network. The simulation results substantiate the complexity analysis results in Section 3.4 that the time complexity of CWA-1L depends heavily on W and the time complexity of CLA-1C depends heavily on the ratio of the lightpath bandwidth capacity (C) to connection granularity (relates to G).

5. Conclusions

Traffic grooming aims at improving resource utilization efficiency in optical networks by effectively packing low-rate connections onto high-rate lightpaths. To achieve this objective, we applied the rerouting approach to the dynamic traffic grooming problem in WDM mesh networks. Two rerouting schemes, rerouting at the lightpath level

(RRLP) and rerouting at the connection level (RRCON), were proposed. As a coarse-granularity rerouting scheme operating at the aggregate lightpath level, RRLP is simple as a rerouting algorithm, but may be less flexible in terms of rerouting choices than RRCON, which is a fine-granularity rerouting scheme operating at the per-flow connection level. Both RRLP and RRCON can achieve a short service disruption time by switching paths after establishing the new path (LSP).

We also proposed the critical-wavelength-avoiding one-lightpath-limited rerouting algorithm (CWA-1L) and the critical-lightpath-avoiding one-connection-limited (CLA-1C) rerouting algorithm based on RRLP and RRCON, respectively. By restricting the rerouting to one lightpath or connection each time for a request, not only are the rerouting algorithms made simpler, but the amount of traffic affected in the rerouting process is also reduced. We compared the performance of CWA-1L and CLA-1C with that of a routing algorithm without the rerouting phase, namely the least virtual hop first algorithm (LVHF). The simulation results show that rerouting reduces the connection blocking probability and therefore improves the resource utilization efficiency. The simulation results also show that CWA-1L generally performs better in relatively dense networks and under relatively light traffic. CLA-1C, on the other hand, is an algorithm with more flexibility and outperforms CWA-1L in relatively sparse networks and under relatively heavy traffic.

Acknowledgments

This work was supported in part by the U.S. National Science Foundation Grants (ANI-0074121 and EPS-0091900) and the UNL Program of Excellence PRISM Priority initiative. An earlier version of this paper was presented at the IEEE Globecom 2004 conference.

References

[1] J. Anderson, J.S. Manchester, A.R. Moral, M. Veeraraghavan, Protocols and architectures for IP optical networking, Bell Labs Technical Journal (1999).
[2] G. Ellinas, E. Bouillet, R. Ramamurthy, J. Labourdette, S. Chaudhuri, K. Bala, Routing and restoration architectures in mesh optical networks, Optical Networks Magazine (2003).
[3] A. Lardies, R. Gupta, R.A. Patterson, Traffic grooming in a multi-layer network, Optical Networks Magazine (2001).
[4] C. Xin, Y. Ye, S. Dixit, C. Qiao, An integrated lightpath provisioning approach in mesh optical networks, in: OFC 2002, March 2002, pp. 547–549.
[5] R. Mahalati, R. Dutta, Reconfiguration of traffic grooming optical networks, in: Proceedings of BroadNets 2004, October 2004, pp. 170–179.
[6] O. Gerstel, P. Lin, G. Sasaki, Combined WDM and SONET network design, in: Proceedings of IEEE INFOCOM'99, NY, vol. 2, March 1999, pp. 734–743.
[7] A.L. Chiu, E. Modiano, Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks, Journal of Lightwave Technology 18 (1) (2000).
[8] J. Wang, V.R. Vemuri, W. Cho, B. Mukherjee, Improved approaches for cost-effective traffic grooming in WDM ring networks: ILP formulations and single-hop and multi-hop connections, IEEE/OSA Journal of Lightwave Technology 19 (2001) 1645–1653.
[9] S. Zhang, B. Ramamurthy, Dynamic traffic grooming algorithms for reconfigurable SONET over WDM networks, IEEE Journal on Selected Areas in Communications 21 (7) (2003) 1165–1172.
[10] K. Zhu, B. Mukherjee, Traffic grooming in an optical WDM mesh network, IEEE Journal on Selected Areas in Communications 20 (1) (2002).
[11] N. Srinivas A1, C. Siva Ram Murthy, Throughput maximization in traffic grooming in WDM mesh networks, Journal of High Speed Networks 13 (2) (2004) 139–154.
[12] W. Yao, B. Ramamurthy, Survivable traffic grooming with path protection at the connection level in WDM mesh networks, IEEE/OSA Journal of Lightwave Technology 23 (10) (2005) 2807–3418.
[13] K. Zhu, B. Mukherjee, On-line approaches for provisioning connections of different bandwidth granularities in WDM mesh networks, in: OFC, March 2002, pp. 549–551.
[14] H. Zhu, H. Zang, K. Zhu, B. Mukherjee, A. novel, generic graph model for traffic grooming in heterogeneous WDM mesh networks, IEEE/ACM Transactions on Networking 11 (2) (2003).
[15] W. Yao, B. Ramamurthy, A link bundled auxiliary graph model for constrained dynamic traffic grooming in optical WDM mesh networks, IEEE Journal on Selected Areas in Communications 23 (8) (2005) 1542–1555.
[16] X. Huang, F. Farahmand, J.P. Jue, Multicast traffic grooming in wavelength-routed WDM mesh networks using dynamically changing light-trees, IEEE Journal of Lightwave Technology 23 (2005) 3178.
[17] K. Lee, V.O.K. Li, A wavelength rerouting algorithm in wide-area all-optical networks, Journal of Lightwave technology 14 (6) (1996).
[18] G. Mohan, C. Siva Ram Murthy, A time optimal wavelength rerouting algorithm for dynamic traffic in WDM networks, Journal of Lightwave technology 17 (3) (1999).
[19] E. Mannie, Generalized multi-protocol label switching (GMPLS) architecture, Internet draft, draft-ietf-ccamp-gmpls-architecture-04.txt, February 2003, work in progress.
[20] H. Harai, M. Murata, H. Miyahara, Performance of alternate routing methods in all-optical switching networks, in: Proceedings of IEEE INFOCOM, Kobe, Japan, vol. 2, April 1997, pp. 516–524.
[21] H.V. Madhyastha, C. Siva Ram Murthy, Efficient dynamic traffic grooming in service-differentiated WDM mesh networks, Computer Networks 45 (2) (2004) 221–235.

Wang Yao is currently a Software R&D Engineer at Blade Network Technologies Inc., Santa Clara, California. He received his Ph.D. degree in the Department of Computer Science and Engineering at the University of Nebraska-Lincoln (UNL) in 2005. He received his B.E. and M.E. degrees from the Beijing University of Aeronautics and Astronautics (China), in 1997 and 2000, respectively. Before he joined UNL, he had worked with Huawei Technologies for one year at Huawei Beijing Research Center as a senior software engineer. His research interests include networking and distributed computing, optical communication networks, and wireless networks.

Byrav Ramamurthy received his B.Tech. degree in Computer Science and Engineering from Indian Institute of Technology, Madras (India) in 1993. He received his M.S. and Ph.D. degrees in Computer Science from University of California (UC), Davis in 1995 and 1998, respectively.
Since August 2003, he has been an associate professor in the Department of Computer Science and Engineering at the University of Nebraska-Lincoln (UNL). He is the Co-director of the UNL Academic Program Priority Initiative in the areas of Simulation & Computing Engineering (SCE) and Information Technology & Telecommunications (ITT). He is the founding co-director of the Advanced Networking and Distributed Experimental Systems (ANDES) Laboratory at UNL. He is the Feature Editor on Theses for the Optical Network Magazine. He served as a guest co-editor for a special issue of IEEE Network magazine on Optical Communication Networks. He served as a member of the technical program committees for the IEEE INFOCOM, IEEE GLOBECOM, Opticomm, ICC and ICCCN confer-

ences. He is author of the textbook "Design of Optical WDM Networks – LAN, MAN and WAN Architectures" published by Kluwer Academic Publishers in 2000. From 2001 to 2003, he served as the founding secretary of the IEEE ComSoc Optical Networking Technical Committee (ONTC).

He was a recipient of the Indian National Talent Search scholarship and was a fellow of the Professors for the Future program at UC Davis. He is a recipient of the UNL Research Council Grant-in-Aid award for 1999, the College of Engineering and Technology Faculty Research Award for 2000 and the UNL CSE Dept. Students Choice Award for the Best Graduate Professor for 2002–2003. His research areas include optical networks, wireless networks, network security,distributed computing, and telecommunications.