Wavelength Rerouting for On-line Multicast in WDM Networks

Yingyu Wan Weifa Liang
Department of Computer Science
The Australian National University
Canberra, ACT 0200, Australia
{yywan,wliang}@cs.anu.edu.au

Abstract

In this paper we consider wavelength rerouting for on-line multicast in all-optical Wavelength Division Multiplexing (WDM) networks where the multicast requests arrive and depart randomly. One limitation of such networks is the wavelength continuity constraint imposed by the all-optical cross-connect switches that requires the same wavelength be used on all the links in a multicast tree. With random arrivals and departures of multicast requests, it happens quite often that a new multicast request has to be blocked due to the fact that there are not enough available resources (e.g. wavelength) to realize the request. Wavelength rerouting, a viable and cost-effective method, has been proposed to improve the blocking probability, which rearranges the wavelengths on certain existing multicast routes to free a wavelength continuous route for the new request. In this paper, we study the wavelength rerouting problem for on-line multicast in both undirected and directed WDM networks with an objective to minimize the disruption incurred to the existing multicast services, or equivalently, to minimize the number of existing multicast routes to be wavelength-rerouted. We first show that the problem is not only NP-hard but also hard to approximate. We then devise approximation algorithms for it with provable approximation guarantees.

1. Introduction

A Wavelength Division Multiplexing (WDM) network consists of optical wavelength routing nodes interconnected by point-to-point fiber links in an arbitrary topology. On each fiber link, the fiber bandwidth is divided into multiple frequency bands (wavelengths) so that several communication requests can be realized over the fiber link at the same time, as long as each uses a different wavelength. There are two types of WDM networks, one allows wavelength conversion at its nodes and the other does not. In a network without wavelength conversion, the realization of a connection request is subject to the *wavelength continuity constraint*, which requires that the same wavelength be used on all the links in the communication path of the request. The wavelength continuity constraint often reduces wavelength utilization because a non-wavelength-continuous route cannot be used for a communication request even it is available. This is especially severe in a network with random arrivals and departures of communication requests. Although wavelength conversion can potentially allow the network to accommodate more communication requests, wavelength converters at nodes are expensive at least in the near future. Hence, most existing work in this area assumes that there is no wavelength conversion in the network. This assumption will be adopted in this paper.

To alleviate the inefficiency due to the wavelength continuity constraint, a viable and cost-effective method called *rerouting* has been proposed as follows. Whenever a new communication request arrives, if there is no wavelength-continuous route for it, rearrange a certain number of existing routes to free a wavelength, in order to make room for the new request. There are two ways to rearrange an existing route. One is "fully rearranging", which is to find a new route with another wavelength to replace the old route. This is also referred to as *nonblocking rearrangement*. Another is "partially rearranging", which keeps the original physical links in the route and reassigns a different wavelength to the route. The latter one is referred to as *wavelength rerouting*. Examples of nonblocking rearrangement and wavelength rerouting can be found in [16, 17, 19, 6] and in [18, 10, 14, 15, 4] respectively.

While rerouting can be used to improve the bandwidth utilization, transmission on each of the rerouted routes must be temporarily shut down to prevent data from being lost during the rerouting process. This causes a low or even zero throughput on those being rerouted traffic. The throughput loss is particularly prominent in all-optical networks wherein each routing traffic is expected to carry gigabits of data flow per second, and hence even a tiny period of outage on a single route will cause significant data loss. Thus,

Proceedings of the 29th Annual IEEE International Conference on Local Computer Networks (LCN'04)
0742-1303/04 $ 20.00 IEEE

minimizing the disruptions (i.e., the number of rerouted routes) is of paramount importance for rerouting in all-optical WDM networks.

It is a general belief that nonblocking rearrangement is much harder than wavelength rerouting because in the former not only a new route for a communication request needs to be found, but also another available wavelength needs to be assigned to each of the links in the new route. Despite the fact that nonblocking rearrangement may improve the blocking probability significantly, compared to wavelength rerouting, it leads to a much longer disruption than expected. In particular in an arbitrary topology network, to make a nonblocking rearrangement for a new multicast request with the minimum disruption is very difficult. In real life, most known works on nonblocking rearrangement are carried out on special topology networks like rings and tori [19].

Multicast is a point-to-multipoint communication that a source node sends a message to multiple destination nodes [21], which has wide application backgrounds including news feeds, video distribution, multimedia conferencing, and so on. To implement a multicast request efficiently, a typical way is to find a tree rooted at the source node and spanning the destination nodes. The multicast message from the source node is then transmitted and propagated to all the destination nodes along the tree links. To facilitate a multicast request in wavelength-routed WDM networks, the concept of *light-tree* and the cross-connect architecture of *splitter-and-delivery* were introduced in [1, 20]. A light-tree is a tree rooted at the source node and spanning the destination nodes, which uses the same wavelength on all its links. Each non-leaf node having more than one child in the tree must have a splitter, which splits the incoming optical signal into multiple identical copies outgoing to its child nodes.

In this paper, we will focus on wavelength rerouting in an arbitrary topology WDM network for on-line multicast. It must be mentioned that, the multicast problem that we will study is totally different from those previous studies on multicast that use different cost metrics [1, 20, 2, 11, 22, 12, 8]. Here, the cost of a new multicast tree is defined as the number of existing multicast trees which have to be rerouted.

1.1. Related Work

Lee and Li [10] first introduced the wavelength rerouting concept. They studied the unicast routing problem with the objective to minimize the disruption incurred due to wavelength rearranging. For an undirected WDM network with n nodes, m physical links and w wavelengths on each link, they proposed a wavelength rerouting scheme called *Parallel Move-To-Vacant Wavelength-Retuning (MTV-WR)*, which has certain advantages. First, it facilitates control because the old and new routes of rerouted requests share the same switching nodes. Second, it reduces the calculation of the optimal rerouting because only the wavelengths of existing routes need to be changed. Third, it significantly reduces the disruption period. An algorithm for implementing the MTV-WR scheme has also been proposed, which takes $O(n^3w + n^2w^2)$ time per unicast request [10]. Mohan and Murthy [15] later provided a $O(n^2w)$ time improved algorithm for the MTV-WR scheme. Caprara et al [4] studied the unicast routing problem in directed WDM networks, which was also referred to as the Venetian Routing (VR) problem. They showed that the problem is not only NP-hard but also hard to approximate. Specifically, when the maximum length of any existing route is no more than 3, the problem is APX-hard, while in general case the problem cannot be approximated within $O(2^{\log^{\frac{1}{2}-\varepsilon} m})$ for any fixed $\varepsilon > 0$ unless $NP \subseteq DTIME(n^{poly \log(n)})$. They also proposed an $O(\sqrt{m/opt})$-approximation algorithm for the problem, where opt is the value of the optimal solution which is between 1 to $n-1$. The main idea of their algorithm is to divide the existing paths into "long" and "short". To find a path for the new request they take all "long" paths and use the "short" paths parsimoniously.

As for non-blocking rearrangement of communications, Saengudomlert et al [19] developed on-line unicast routing and wavelength assignment algorithms for bidirectional rings and tori of n nodes with the objective to minimize the number of wavelengths to support given traffic matrices dynamically. For a bidirectional ring, their algorithm requires at most three lightpath rearrangements per unicast request. For a torus, their algorithm requires at most $\sqrt{n}-1$ lightpath rearrangements per unicast request. Chen and Modiano [6] studied the off-line unicast problem in a reconfigurable bidirectional ring in which wavelength converters at nodes are installed, with an objective to minimize the number of wavelengths. Non-blocking rearrangement has also been used to improve the capacity utilization of survivable routing in WDM networks [13, 3]. For example, Liu et al [13] proposed an algorithm for rerouting the backup path successively to reduce the redundancy for each primary path, while Bouillet et al [3] proposed an algorithm for rerouting both primary and backup paths simultaneously.

1.2. Contributions

Inspired by the works of Lee and Li [10] and Caprara et al [4], in this paper we consider the Wavelength Rerouting problem for on-line Multicast (WRM) in both undirected and directed WDM networks. Our major contributions are as follows.

If the considered WDM network is undirected, we first show that WRM is not only NP-hard but also hard to approximate. In other words, the problem cannot be ap-

Wavelength Rerouting Scheme for On-Line Multicast

1. For each wavelength $\lambda_i \in \Lambda$ do

1.1 For each light-tree $T \in \mathcal{T}_i$, if there is another available wavelength $\lambda_j \in \Lambda$ ($i \neq j$) on each link in T, i.e., $\forall e \in E(T)$, $\lambda_j \in \Lambda_e$, then T can migrate to λ_j from λ_i, and T is said to be *tunable*. Otherwise, T is said to be *untunable*. Let $\mathcal{T}_i'$ be the subset of $\mathcal{T}_i$ which contains all the tunable light-tree, and $\mathcal{T}_i'' = \mathcal{T}_i - \mathcal{T}_i'$.

1.2 Find a light-tree T' rooted at s and spanning the nodes in D such that
(i) wavelength λ_i is assigned to each of its links;
(ii) the number of light-trees in $\mathcal{T}_i'$ link-intersecting with T' is minimized;
(iii) there is no any link-intersection between T' and any light-tree in $\mathcal{T}_i''$.

1.3 If there is such a T' in the network, denote it by $T_{s,d}^i$. The cost of $T_{s,d}^i$ is defined as the number of intersected light-trees in $\mathcal{T}_i'$. Otherwise, let $T_{s,d}^i = \emptyset$ and its cost be ∞.

2. Select a light-tree with the minimum cost from at most w trees, $T_{s,d}^1, T_{s,d}^2, \cdots$, and $T_{s,d}^w$. Assume that $T_{s,d}^i$ has been chosen. If the cost of $T_{s,d}^i$ is ∞, then the multicast request $(s; D)$ cannot be supported by the wavelength rerouting scheme and will be rejected. Otherwise, each light-tree in $\mathcal{T}_i'$ that link-intersects with $T_{s,d}^i$ is migrated to another available wavelength. As a result, a route for the multicast request $(s; D)$ is found, which is $T_{s,d}^i$ with wavelength λ_i.

Figure 1. Overview of the wavelength rerouting scheme

proximated within $(1-\varepsilon)\ln|D|$ for any constant ε unless $NP \subset DTIME(n^{\log\log n})$, where D is the set of destination nodes and $0 < \varepsilon < 1$. This inapproximability result still holds even if the network is planar. Instead, we then present an approximation algorithm for it which delivers a solution within $O(\log|D|)$ times the optimum. If the network is directed, following a result in [4], WRM is inapproximable within $O(2^{\log^{\frac{1}{2}-\varepsilon} m})$ for any constant $\varepsilon > 0$ unless $NP \subseteq DTIME(n^{poly\log(n)})$ due to the problem they dealt with is a special case of WRM. Using the similar techniques of [4], we are able to devise an $O(\sqrt{m \cdot \alpha/opt})$-approximation algorithm for it, where *opt* is the number of existing multicast routes to be rerouted in an optimal solution of WRM, $1 \leq opt \leq n-1$, and α is the best known approximation guarantee for the directed Steiner tree problem. Currently $\alpha \leq i(i-1)|D|^{1/i}$ for any fixed $i > 1$ following the result in [5].

The rest of this paper is organized as follows. In Section 2 we define the network model. In Section 3 we provide an overview of the wavelength rerouting scheme for on-line multicast. In Sections 4 and 5 we focus on studying WRM in undirected and directed WDM networks respectively. We conclude our discussions in Section 6.

2. Network Model

A WDM network $M = (N, L)$ without wavelength conversion is considered, where N is the set of communication nodes and L is the set of fiber links. The bandwidth of each link is divided into a set Λ of wavelengths, where $\Lambda = \{\lambda_1, \lambda_2, \cdots, \lambda_w\}$. It is assumed that each node has sufficient optical transmitters, receivers and splitters to realize the routes so that no new communication request will be blocked due to the lack of these resources.

A multicast request is denoted by a pair $(s; D)$, where s is the source node and D is the set of destination nodes. We assume that a new multicast request arrives and departs randomly. Denote by Λ_e ($\subseteq \Lambda$) the set of available wavelengths on link e at the time a multicast request arrives. Let $\mathcal{T}$ be the set of light-trees and each light-tree corresponds to an existing route for an existing multicast request. $\mathcal{T}$ can be further partitioned into w disjoint subsets $\mathcal{T}_1, \mathcal{T}_2, \cdots, \mathcal{T}_w$, according to which wavelength is used in the trees in the set. Obviously, any two light-trees in the same subset $\mathcal{T}_i$ for $1 \leq i \leq w$ are pairwise link-disjoint, i.e., any two light-trees in the set do not share any link.

3. Overview of the Wavelength Rerouting Scheme

To accommodate a new multicast request $(s; D)$, we need to find a light-tree rooted at s and spanning the destination nodes such that there is a wavelength available on all the links in the tree. This can be done within the following two phases.

1) Find a light-tree without rerouting any existing multicast routes.

2) Find a light-tree with rerouting some existing routes if Phase 1) fails.

Phase 1) can be easily implemented as follows. For each wavelength $\lambda \in \Lambda$, check whether there is a connected subgraph induced by an available wavelength λ such that node s and the nodes in D are included in it. If so, a light-tree assigned with wavelength λ for the multicast request is obtained. Otherwise, a wavelength rerouting scheme similar to MTV-WR [10] will be adopted, which is described in Fig. 1.

In Step 1.2 of the scheme, the requirement of link-intersecting with the minimum number of light-trees in $\mathcal{T}_i'$ is to minimize the disruption for rerouting. Steps 1.1, 1.3 and 2 are easy to implement. To implement Step 1.2, we construct a graph $G = (V, E)$ where $V = N$ and $E = \{e \in L \mid \lambda_i \in \Lambda_e$ or $\exists T \in \mathcal{T}_i', e \in E(T)\}$. The rest is to find a multicast tree in G rooted at node s and spanning the nodes in D which link-intersects the minimum number of light-trees in $\mathcal{T}_i'$. This leads to a combinatorial optimization problem which is defined as follows.

Definition 1 (Wavelength Rerouting for Multicast, WRM) *Given a graph $G = (V, E)$ with $n = |V|$ nodes and $m = |E|$ edges, a collection $\mathcal{T}$ of multicast trees which are pairwise link-disjoint, a multicast request $(s; D)$, $s \in V$ and $D \subseteq V \setminus \{s\}$, the objective is to construct a multicast tree rooted at s and spanning the nodes in D such that the number of multicast trees in $\mathcal{T}$ that link-intersects with the multicast tree is minimized.*

In what follows we focus on studying the computational complexity of WRM and devising approximation algorithms for it in both undirected and directed WDM networks.

4. Undirected WDM Networks

4.1. Hardness of Approximation

We show that WRM in undirected WDM networks is not only NP-hard but also hard to approximate through a reduction from the Set Cover (SC) problem. Recall that an instance of SC is given by a finite set S and a collection $\mathcal{C}$ of subsets of S. The objective is to find a minimum set cover, i.e., a minimum-cardinality collection $\mathcal{C}'$ of subsets in $\mathcal{C}$ such that each element in S is contained by at least a set in $\mathcal{C}'$. Feige [7] has shown that SC is inapproximable within $(1 - \epsilon) \ln |S|$ for any constant ϵ with $0 < \epsilon < 1$ unless $NP \subset DTIME(n^{\log \log n})$.

Theorem 1 *Given a planar, undirected graph $G(V, E)$ and a multicast request $(s; D)$, WRM cannot be approximated within $(1 - \epsilon) \ln |D|$ for any constant ϵ with $0 < \epsilon < 1$ unless $NP \subset DTIME(n^{\log \log n})$.*

Proof Given an instance I of SC, with $S = \{d_1, d_2, \cdots, d_k\}$ and $\mathcal{C} = \{C_1, C_2, \cdots, C_l\}$, a planar graph $G = (V, E)$ is constructed as follows, shown in Fig. 2.

We first construct an undirected graph $G' = (V', E')$ from I. The nodes in G' are partitioned into three disjoint sets which correspond to three layers and the edges exist

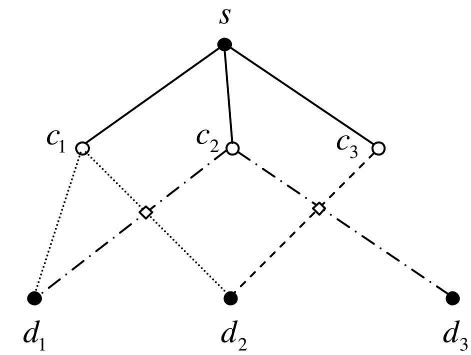

Figure 2. Reduction from SC to WRM

only between the nodes at adjacent layers. In the first layer there is only a source node s. In the second layer each node is a *set node* c_i for each subset C_i in $\mathcal{C}$, which is represented by a white round node. There is an edge (s, c_i) for each set node c_i . In the third layer there is an *element node* d_i corresponding to each element d_i in S, which is represented by a black round node. There is an edge (c_i, d_j) in E' if $d_j \in C_i$. Then, for each set node c_i, construct a multicast tree T_i including c_i and spanning its neighboring element nodes. Fig.2 illustrates three multicast trees, represented by dotted, dashed, and dash-dot lines.

We then embed G' into a plane to obtain a planar graph $G = (V, E)$ as follows. The nodes in G' are distributed into three layers. In the embedding if there are two edges to cross each other, create a new *crossing node* at the crossing point and break the two original edges into four new edges. The crossing nodes are represented by white diamond nodes. We repeat this procedure until the resulting graph is planar. As a result, a planar graph is obtained. Fig.2 illustrates the construction of a planar graph for an instance of SC, where $S = \{d_1, d_2, d_3\}$ and $\mathcal{C} = \{C_1 = \{d_1, d_2\}, C_2 = \{d_1, d_3\}, C_3 = \{d_2\}\}$. T_1 is represented by the dotted lines, T_2 is represented by the dash-dot lines, and T_3 is represented by the dashed lines.

Let $(s; D)$ be a multicast request and $D = \{d_1, d_2, \cdots, d_k\}$. Thereafter, we get an instance I_{WRM} of WRM. The transformation is clearly polynomial. Now, we show that there is a set cover with cardinality q in I if and only if there is a multicast tree spanning the nodes in $\{s\} \cup D$ that link-intersects with at most q existing multicast trees in I_{WRM}.

Assume that there is a set cover $\mathcal{C}'$ such that $|\mathcal{C}'| = q$. Without loss of generality, let $\mathcal{C}' = \{C_1, C_2, \cdots, C_q\}$. We construct a multicast tree T as follows. Add the edge (s, c_i) to T for $1 \leq i \leq q$. Add all the edges in the multicast tree T_i to T for $1 \leq i \leq q$. For each node $d_j \in D$, if $d_j \in C_i$ in I, there is a path in T from s to d_j which consists of edge (s, c_i) and the path from c_i to d_j in the multicast tree T_i.

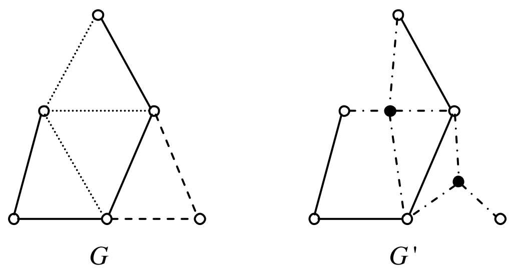

Figure 3. Construction of the auxiliary graph G' from G

So, T is a multicast tree for the multicast request $(s; D)$ that link-intersects with exactly q existing multicast trees.

If there is a multicast tree T for the multicast request $(s; D)$ and link-intersecting only q existing multicast trees, then a set cover is constructed as follows. For each node $d_j \in D$, there is at least one edge in T incident to it. If there are more than one edge incident to it, select one from these edges arbitrarily. If the edge is in the multicast tree T_i, then T_i is link-intersecting with T. From the above construction, the element d_j is contained in the subset C_i in I. Add the set C_i to the set cover. Thus, a set cover for I is obtained and its cardinality is at most q.

The above reduction is cost-preserving, which means that for any set cover $\mathcal{C}'$, its corresponding multicast tree T has the same cost as $\mathcal{C}'$, or *vice versa*. Thus, if there is a polynomial time approximation algorithm for WRM with an approximation guarantee better than $(1-\epsilon)\ln|D|$ for a constant $0 < \epsilon < 1$, there must be a polynomial time $(1-\epsilon)\ln|S|$-approximation algorithm for SC. The theorem follows the result given by Feige [7] that SC cannot be approximated within $(1-\epsilon)\ln|S|$ for any constant $0 < \epsilon < 1$ unless $NP \subset DTIME(n^{\log\log n})$. $\Box$

4.2. Approximation Algorithm

We now present an approximation algorithm for WRM. The basic idea behind the proposed algorithm is to construct an auxiliary graph $G' = (V', E')$ from $G(V, E)$ and to transform WRM in G into a minimum node-weighted Steiner tree problem in G'.

Given an undirected graph $G = (V, E)$ and a collection $\mathcal{T} = \{T_1, T_2, \cdots, T_k\}$ of multicast trees which are pairwise link-disjoint, a node-weighted auxiliary graph $G' = (V', E')$ is constructed as follows. Initially, let $G' = G$. For each multicast tree T_i $(1 \le i \le k)$, remove the edges in T_i from G' and add a *tree node* t_i to G' at the same time, which is the representation of T_i. For each node $v \in V$, add a new edge (t_i, v) to G' if $v \in V(T_i)$, where $V(T_i)$ is the set of nodes in T_i. Fig.3 illustrates the construction of G'. There are two multicast trees in G, the tree edges are represented by the dotted lines and dashed lines respectively. Two black nodes in G' represent the two tree nodes and the dash-dot edges represent the newly added edges. Assign weights to the nodes in G' as follows. Assign 1 to each tree node and 0 to each of the other nodes. Find a Steiner tree T' in G' spanning the nodes in $\{s\} \cup D$ to minimize the weighted sum of the nodes in T'. A multicast tree T for the multicast request $(s; D)$ in G is then obtained by replacing each tree node in T' with the edges of the multicast tree represented by the tree node. Obviously, the weighted sum of the nodes in T' is equal to the cost of T. On the other hand, let opt_G be the cost of an optimal multicast tree in G and $opt_{G'}$ be the weighted sum of the nodes in a minimum node-weighted Steiner tree in G'. Following the construction of G', we have $opt_{G'} \le opt_G$. Thus, an approximate, minimum node-weighted Steiner tree can be found, whose cost is $O(\log|D|)$ times the optimum, by applying the approximation algorithm for the node-weighted Steiner tree problem in [9]. We thus have the following theorem.

Theorem 2 *In undirected WDM networks, there is an approximation algorithm for WRM which delivers a solution within $O(\log|D|)$ times the optimum.*

Following the result given in Theorem 1, the proposed algorithm is almost optimal in terms of the approximation guarantee.

5. Directed WDM Networks

It is a general belief that WRM in directed networks is much harder than it in undirected networks, because the multicast tree must be a directed tree. Therefore, the technique to solve the problem in undirected networks cannot be applied to solve it in directed networks. Caprara et al [4] have already shown that wavelength rerouting for on-line unicast in directed networks is not only NP-hard but also hard to approximate. While it is well known that the unicast problem is a special case of the multicast problem, the problem considered here is at least as hard as the unicast problem. In other words, WRM in directed networks cannot be approximated within $O(2^{\log^{\frac{1}{2}-\epsilon} m})$ for any fixed $\epsilon > 0$ unless $NP \subseteq DTIME(n^{poly\log(n)})$. Thus, in the following we focus on devising an approximation algorithm for it.

5.1. Approximation Algorithm

We use the similar technique in [4] in the design of the proposed algorithm. That is, we partition the existing multicast trees into "large" and "small" trees. To find a new multicast tree, we take all "large" trees and we use the "small" trees parsimoniously. Note that there cannot be too many

Algorithm `WRM-DB`(G,$\mathcal{T}$,s,D,B)
Input: The directed graph G, the set $\mathcal{T}$ of existing multicast trees, a new multicast request $(s; D)$, and a constant B which is an upper bound on the number of nodes in each existing multicast tree in $\mathcal{T}$.
Output: T_{WRM}, a multicast tree for $(s; D)$.
Begin
1. Construct an auxiliary edge-weighted directed graph G' from G.
2. Find a directed Steiner tree T_{DST} in G' rooted at s and spanning nodes in D, by applying the approximation algorithm in [5].
3. Transform T_{DST} into a multicast tree T_{WRM} for the multicast request $(s; D)$ in G.
End

Figure 4. Approximation algorithm for WRM in directed graphs where the number of nodes of each existing multicast tree is bounded by a constant B

of such "large" trees. In the design of their constant factor approximation algorithm for "short" paths, the original graph is transformed to an auxiliary graph, which cannot be applied to multicast routing. We design a different approach to construct an auxiliary graph for multicast routing. In what follows, we first consider a special case of the problem where the number of nodes in each existing multicast tree is bounded by a constant B, for which we present an approximation algorithm in Fig 4.

We now give the implementation details of algorithm `WRM-DB`. The construction of G' is as follows. Each existing multicast tree $T \in \mathcal{T}$ is transformed into another corresponding graph G_T. Without loss of generality, let $\{v_1, v_2, \cdots, v_l\}$ be the set of child nodes of u for each non-leaf node u in T. Remove the edges $< u, v_1 >, < u, v_2 >, \cdots$ and $< u, v_l >$ from T, add a new node u_T and a new edge $< u, u_T >$ to T, and assign the new edge with weight 1. Add a new edge $< u_T, v_i >$ to T and assign it with weight 0 for $1 \leq i \leq l$. If v is a child of u and is a non-leaf node of T, then a new node v_T has been added to the tree in the construction. Now, add a new edge $< u_T, v_T >$ and assign it with weight 0 to T. Denote by G_T the resulting graph, which is shown in Fig. 5(b). Note that G_T is a directed acyclic graph (DAG).

The graph G' is then obtained from G by replacing each existing multicast tree T with G_T. Assign weight 0 to each of the other edges that are not in any existing multicast tree.

The directed Steiner tree T_{DST} in G' can be transformed into a multicast tree T_{WRM} in G by reversing the above transformation, which is described as follows. Initially, define an auxiliary directed graph $G_{WRM} = (V, \emptyset)$. For each existing multicast tree T, if T_{DST} contains at least an edge of the subgraph G_T, add all the edges in T to G_{WRM}. For the other edges in T_{DST} that are not in any subgraph G_T, add them to G_{WRM}. The resulting graph G_{WRM} is a connected subgraph of G including node s and the nodes in D.

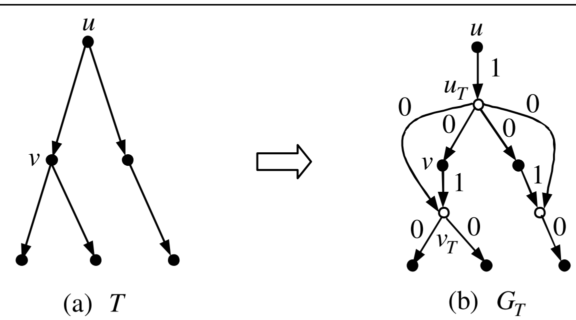

Figure 5. Construction of G_T from T

Note that there must be a directed path in G_{WRM} from s to d for each destination node $d \in D$ because there is a directed path in T_{DST} from s to d. A multicast tree T_{WRM} is finally derived from G_{WRM} by removing some edges from the graph.

We now bound the cost of T_{WRM}. Let T_{OPT} be an optimal multicast tree for the multicast request $(s; D)$ in G and *opt* the cost of it. Clearly, *opt* is between 1 and $n - 1$ because a multicast tree contains at most $n - 1$ edges. Thus, T_{OPT} intersects with at most $n - 1$ existing multicast trees. Let opt_{DST} be the weighted sum of the edges in a minimum directed Steiner tree in G' for the multicast request $(s; D)$. Let $c(T_{WRM})$ be the cost of T_{WRM} and $c(T_{DST})$ the weighted sum of the edges in T_{DST}. Let $|V(T)|$ be the number of nodes of an existing multicast tree T. We have the following lemma.

Lemma 1 $opt_{DST} \leq \lfloor B/2 \rfloor \cdot opt$.

Proof For any given multicast tree T in G, it has been transformed into a subgraph G_T in G'. We show that there is a tree in G' which spans all the nodes in T and the weighted sum of the edges in T is no more than $\lfloor |V(T)|/2 \rfloor$.

Algorithm `WRM-D`(G,$\mathcal{T}$, s, D)
Input: The given directed graph G, the set $\mathcal{T}$ of existing multicast trees,
and a new multicast request $(s; D)$.
Output: T_{WRM}, a multicast tree for the multicast request $(s; D)$.
Begin
1. For each $2 \le x \le n-1$ do
1.1 $\mathcal{Q}_x = \mathcal{T}$;
1.2 For each existing multicast tree $T \in \mathcal{T}$, if T contains more than x nodes,
remove it from $\mathcal{Q}_x$, i.e., $\mathcal{Q}_x = \mathcal{Q}_x \setminus \{T\}$;
1.3 Find a multicast tree for the multicast request $(s; D)$ by calling `WRM-DB`(G,$\mathcal{Q}_x$,s, D,x).
Denote by T^x_{WRM} the resulting tree;
2. Choose a multicast tree T_{WRM} with the minimum cost from $\{T^2_{WRM}, T^3_{WRM}, \cdots, T^{n-1}_{WRM}\}$.
End

Figure 6. Approximation algorithm for WRM in directed graphs

Let $r(T)$ be the root of T. Denote by $level(u)$ the distance from $r(T)$ to node u in T for each u. Clearly, $level(r(T)) = 0$. The tree is traversed by starting from the root using the Breadth-First-Search (BFS) technique. A queue Q is used to keep the trace of the BFS procedure. Initially $Q = \{r(T)\}$. While Q is not empty, proceed as follows. Remove a node u from the head of Q. Visit u, u_T and edge $< u, u_T >$ in G_T. For each child node v of u in T, visit node v and edge $< u_T, v >$ in G_T. If v is a non-leaf node in T, then visit node v_T and edge $< u_T, v_T >$ in G_T. For each child node w of v in T, visit node w and edge $< v_T, w >$, and insert w into the tail of Q. Repeat the BFS procedure until Q is empty. As a result, all the visited nodes and edges compose a tree in G' which spans the nodes of T. In the procedure, only when a node u is a non-leaf node in T and $level(u)$ is even, the edge $< u, u_T >$ with weight 1 is visited. So the weighted sum of the edges in the resulting tree is no more than $\lfloor |V(T)|/2 \rfloor$.

Assume that T is an existing multicast tree link-intersecting with T_{OPT}. Let $E(T, T_{OPT})$ be the set of edges in T that intersects with the edges in T_{OPT} and $C^T_1, C^T_2, \cdots$, and C^T_j the connected components of T induced by the edges in $E(T, T_{OPT})$. Assume that n_i is the number of nodes in C^T_i, $1 \le i \le j$. Following the construction, a directed Steiner tree A_{DST} in G' that corresponds to T_{OPT} can be obtained. Let $C'_{T_1}, C'_{T_2}, \cdots$, and C'_{T_j} be the corresponding connected components of A_{DST} in G' induced from $C^T_1, C^T_2, \cdots, C^T_j$ in G. Note that each of these connected components is a tree, and the weighted sum of the edges in C'_{T_i} is at most $\lfloor n_i/2 \rfloor$, for $1 \le i \le j$. Thus, the contribution of T to the weighted sum of the edges in tree A_{DST} is at most

$$\sum_{1 \le i \le j} \lfloor \frac{n_i}{2} \rfloor \le \lfloor \frac{\sum_{1 \le i \le j} n_i}{2} \rfloor \le \lfloor \frac{B}{2} \rfloor.$$

Since $opt_{DST} \le c(A_{DST})$, the lemma then follows. □

Lemma 2 *Algorithm* `WRM-DB` *is* $\lfloor B/2 \rfloor \alpha$*-approximation, where* α *is the best known approximation guarantee for the directed Steiner tree problem.*

Proof Following the above discussion, we have $c(T_{WRM}) \le c(T_{DST})$ and $c(T_{DST}) \le \alpha \cdot opt_{DST}$. From Lemma 1, $opt_{DST} \le \lfloor B/2 \rfloor \cdot opt$. So algorithm `WRM-DB` is $\lfloor B/2 \rfloor \alpha$-approximation. □

An approximation algorithm for WRM is given in Fig. 6. For each nonnegative integer x with $2 \le x \le n-1$, define $f(G, x)$ as the number of existing multicast trees in G with each containing more than x nodes.

Lemma 3 $f(G, x) \le m/x$ *and* $c(T^x_{WRM}) \le \alpha \cdot \lfloor x/2 \rfloor \cdot opt + f(G, x)$.

Proof Since all existing multicast trees in $\mathcal{T}$ are pairwise link-disjoint, the number of multicast trees containing more than x nodes in the network is at most $\lfloor m/x \rfloor$. So, $f(G, x) \le m/x$.

From Lemma 2, the multicast tree T^x_{WRM} returned by `WRM-DB`(G,$\mathcal{Q}_x$,s,D,x) link-intersects at most $\alpha \cdot \lfloor x/2 \rfloor \cdot opt$ multicast trees in $\mathcal{T}_x$. In the worst case T^x_{WRM} link-intersects all the multicast trees in $\mathcal{T}$ that have more than x nodes. Thus $c(T^x_{WRM}) \le \alpha \cdot \lfloor x/2 \rfloor \cdot opt + f(G, x)$. □

T_{WRM} is the one with the minimum cost among the trees $T^2_{WRM}, T^3_{WRM}, \cdots$, and T^{n-1}_{WRM}. So,

$$\begin{aligned} c(T_{WRM}) &= \min_{2 \le x \le n-1} \{c(T^x_{WRM})\} \\ &\le \min_{2 \le x \le n-1} \{\alpha \cdot \lfloor x/2 \rfloor \cdot opt + f(G, x)\}, \end{aligned}$$

which is minimized when x is equal to $\lceil \sqrt{2 \cdot m/(\alpha \cdot opt)} \rceil$ approximately. Thus,

$$c(T_{WRM}) \le \sqrt{2 \cdot m \cdot \alpha \cdot opt}.$$

Theorem 3 *In directed WDM networks, there is a polynomial-time approximation algorithm for WRM, which delivers a solution within $O(\sqrt{m \cdot \alpha / opt})$ times the optimum, where opt is the cost of an optimal solution with* $1 \le opt \le n-1$.

6. Conclusions

In this paper we studied the wavelength rerouting for on-line multicast (WRM) in both undirected and directed WDM networks. We showed that WRM is not only NP-hard but also hard to approximate. Instead, we devised polynomial-time approximation algorithms for the problem with provable approximation guarantees in both networks.

Acknowledgment

It is acknowledged that the work by the authors was supported by research grant #DP0449431, funded by Australian Research Council under its Discovery Schemes.

References

[1] M. Ali and J. S. Deogun. Power-efficient design of multicast wavelength-routed networks. *IEEE Journal on Selected Areas in Communications*, 18(10):1852–1862, 2000.

[2] J. C. Bermond, L. Gargano, S. Perennes, A. Rescigno, and U. Vaccaro. Efficient collective communication in optical networks. *Theoretical Computer Science*, 233(1-2):165–189, 2000.

[3] E. Bouillet, J. F. Labourdette, R. Ramamurth, and S. Chaudhuri. Lightpath re-optimization in mesh optical networks. *7th European Conference on Network and Optical Networks*, 2002.

[4] A. Caprara, G. F. Italiano, G. Mohan, A. Panconesi, and A. Srinivasan. Wavelength rerouting in optical networks, or the venetian routing problem. *Journal of Algorithms*, 45(2):93–125, 2002.

[5] M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed steiner problems. *Journal of Algorithms*, 33(1):73–91, 1999.

[6] L.-W. Chen and E. Modiano. Efficient routing and wavelength assignment for reconfigurable wdm networks with wavelength converters. In *Proc. IEEE INFOCOM'03*, pages 1785 – 1794, San Francisco, CA, April 2003.

[7] U. Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45:634–652, 1998.

[8] X. D. Hu, T. P. Shuai, X. Jia, and M. H. Zhang. Multicast routing and wavelength assignment in wdm networks with limited drop-offs. In *Proc. IEEE INFOCOM'04*, Hong Kong, China, March 2004.

[9] P. Klein and R. Ravi. A nearly best-possible approximation algorithm for node-weighted steiner trees. *Journal of Algorithms*, 19(1):104–115, 1995.

[10] K. C. Lee and V. O. K. Li. A wavelength rerouting algorithm in wide-area all-optical networks. *Journal of Lightwave Technology*, 14:1218–1229, 1996.

[11] W. Liang and H. Shen. Multicasting and broadcasting in large wdm networks. In *Proc. IEEE IPPS'98*, pages 516–523, Orlando, FL, 1998.

[12] R. Libeskind-Hadas and R. Melhem. Multicast routing and wavelength assigment in multihop optical networks. *IEEE/ACM Transactions on Networking*, 10(5):621–629, 2002.

[13] Y. Liu, D. Tipper, and P. Siripongwutikorn. Approximating optimal spare capacity allocation by successive survivable routing. In *Proc. IEEE INFOCOM'01*, pages 699–708, Anchorage, AK, April 2001.

[14] G. Mohan and C. Murthy. Efficient algorithms for wavelength rerouting in wdm multi-fiber unidirectional ring networks. *Computer Communication*, 22:232–243, 1999.

[15] G. Mohan and C. Murthy. A time optimal wavelength rerouting for dynamic traffic in wdm networks. *Journal of Lightwave Technology*, 17:406–417, 1999.

[16] A. Narula-Tam, P. J. Lin, and E. Modiano. Efficient routing and wavelength assignment for reconfigurable wdm networks. *IEEE Journal on Selected Areas of Communications*, 20:75–88, 2002.

[17] P. Saengudomlert, E. Modiano, and R. Gallager. An on-line routing and wavelength assignment algorithm for dynamic traffic in a wdm bidirectional ring. In *Proc. Joint Conference on Information Sciences*, pages 1331–1334, Durham, NC, March 2002.

[18] P. Saengudomlert, E. Modiano, and R. Gallager. Dynamic wavelength assignment for wdm all optical tree networks. *Allerton Conference on Communications, Control and Computing*, October 2003.

[19] P. Saengudomlert, E. Modiano, and R. Gallager. On-line routing and wavelength assignment for dynamic traffic in wdm ring and torus networks. In *Proc. IEEE INFOCOM'03*, pages 1805 – 1815, San Francisco, CA, April 2003.

[20] L. H. Sahasrabuddhe and B. Mukherjee. Light-trees: optical multicasting for improved performance in wavelength-routed networks. *IEEE Communications Magazine*, 37(2):67–73, 1999.

[21] L. H. Sahasrabuddhe and B. Mukherjee. Multicast routing algorithms and protocols: a tutorial. *IEEE Network*, 14(1):90–102, 2000.

[22] X. Zhang, J. Wei, and C. Qiao. Constrained multicast routing in wdm networks with sparse light splitting. In *Proc. IEEE INFOCOM'00*, pages 1781–1790, Tel Aviv, Israel, March 2000.