Online Multicasting in WDM Networks with Shared Light Splitter Bank

Yuzhen Liu and Weifa Liang

Department of Computer Science, The Australian National University Canberra, ACT 0200, Australia {yliu,wliang}@cs.anu.edu.au

Abstract. We study online multicasting in WDM networks with shared light splitter bank. Our objective is to maximize the network throughput. It is desirable that the cost of realizing each multicast request be minimized, and the network throughput will be maximized ultimately through the cost saving on each individual request. We first propose a cost model for realizing an online multicast request under such network environments with limited light splitters and wavelength converters, which models the cost of utilization of network resources, particularly in modelling the light splitting and wavelength conversion ability at nodes. We then show that finding a cost-optimal multicast tree for a multicast request under the proposed cost model is NP-complete, and instead devise approximation and heuristic algorithms for it. We finally conduct experiments to evaluate the performance of the proposed algorithms. The results show that the proposed algorithms are efficient and effective in terms of network throughput.

1 Introduction

A WDM network consists of optical wavelength routing nodes interconnected by point-to-point fiber links. In WDM networks, the fiber bandwidth is partitioned into multiple frequency bands (wavelengths) that may be used to transmit messages simultaneously as long as each message uses a different wavelength. It becomes increasingly evident that WDM networks are able to provide data transmission rates several orders of magnitude higher than current electronic networks, and will soon become the core technology for the next generation Internet by providing unprecedented large available wavelengths [1]. The key to high speed in these networks is to maintain signals in optical form rather than traditional electronic form with the devices such as optical crossconnects(OXCs) [2]. The OXCs equipped with light splitters, referred to as multicast-capable OXCs (MC-OXCs), have the ability to split the incoming signal into more than one outgoing signal with the same wavelength but lower power level. The splitters are the fundamental optical device contributing to power loss [3]. Even in the ideal case the power of each output of a splitter is only 1/n of that of the input signal, where n is the fanout of the splitter. Some devices such as erbium-doped fiber amplifiers can be used to keep the power level of an optical signal above some threshold so that the signal is able to be detected. However, the use of amplifiers would increase the cost of WDM networks. Another factor increasing the cost of WDM networks is the use of wavelength converters that allow the wavelength of outgoing signals different from that of their incoming counterpart. Therefore, the number of light splitters and the number of wavelength converters in a network should be taken into account [3,4], and the splitters and converters installed at a node should be shared by all its incoming signals in a power-efficient and cost-effective WDM network.

Light splitters are the key components to implement multicast. Multicast routing and wavelength assignment (MC-RWA) is a fundamental problem in WDM networks, which aims at finding a set of links and wavelengths on these links to establish the connection from the source to the destination nodes. MC-RWA includes the building of a routing tree (light-tree) and the assignment of wavelengths to the links in the tree. Since the combined multicast routing and wavelength assignment is a hard problem, the most adopted strategy is to decouple the problem into two separate subproblems: the light-tree routing problem and the wavelength assignment problem [2,3,5]. The former aims to build a routing tree for a multicast request, while the latter aims to assign the available wavelengths to the links in the tree.

There are several studies for constructing multicast trees under some physical constraints on optical networks. Sahin and Azizolgu [6] considered the multicast problem under various fanout policies. Libeskind-Hadas and Melhem [7] investigated multicast communication in circuit-switched multi-hop networks by showing, despite the fact that the general multicast problem is NP-hard, it is polynomially solvable when the optimization objective is the wavelength assignment only. For constructing constrained multicast trees in WDM networks, Bermond et al [8] investigated routing and wavelength assignment with only unicast-capable switches. Libeskind-Hadas [9] extended the unicast communication (point-to-point communication model) by proposing a multi-path routing model, in which the multicast problem is to find a set of paths from the source to the destination nodes such that each path contains a subset of destination nodes, all the destination nodes are included by these paths, and the cost sum of these paths is minimized. Zhang et al [10] considered the multicast routing problem by focusing on the limited light splitting of optical switches and provided several heuristics for the problem. Zhang and Yang [5] considered the multicast problem with an objective of minimizing the number of converters by providing an approximation algorithm with approximation ratio of $O(\log n)$. Rouskas [2] provided an excellent survey of the problem under the light splitter switching model.

In this paper we consider the online multicast routing and wavelength assignment problem in the networks in which light splitters/wavelength converters are installed only on a fraction of the nodes and shared by incoming signals. Since there are efficient algorithms for wavelength assignment in tree structures available [5,11], we focus on the routing problem under this shared light splitter bank architecture. Specifically we consider the following online mulicasting problem.

Assume that there is a sequence of multicast requests that is unknown in advance and the requests arrive one by one. Once a multicast request arrives, the response by the system to the request is to either realize the request by building a multicast tree for it or reject the request due to lack of network resources. The objective is to maximize the network throughput that is the number of the realized multicast requests in the given sequence. Due to the unknown pattern of future requests, we focus on realizing each individual multicast request by building an economic multicast tree for the request.

2 Preliminaries

In this section we first introduce a model of WDM networks with shared light splitter bank. We then propose a node cost model that characterizes the utilization of these network resources. We finally define the optimal multicast tree problem and the online multicast request maximization problem.

2.1 Shared Light Splitter Bank Model

The WDM network with shared light splitter bank is modelled by an *undirected* graph $G(V, E, \Lambda, w)$, where V is a set of nodes (vertices), E is a set of bidirectional optical fiber links (edges), $\Lambda = \{\lambda_1, \lambda_2, \dots, \lambda_k\}$ is a set of wavelengths in G and $\Lambda(e)$ is the set of wavelengths available on edge e, and w is a function from V to the set of non-negative real number, |V| = n, |E| = m, and |A| = k. The node weight w(v) models the light splitting/wavelength conversion ability at v and is inversely proportional to this ability, that is, a node with higher ability has a smaller node weight, whereas a node with lower ability has a larger node weight. For example, w can be such a function defined as follows: w(v) = 1 - f(v)if $f(v) \neq 0$; otherwise $w(v) = \infty$, where f(v) is the ratio of the residual light splitting/wavelength conversion ability to the initial light splitting/wavelength conversion capacity at v with $0 \le f(v) \le 1$. When f(v) = 0, all messages entering v will be trapped at v and there is no outgoing flow from v. Thus, we may set w(v) to be a sufficiently large positive value, and v is unlikely to be included as an internal node into multicast trees. When f(v) = 1, which means that there is no traffic load at v at all or the light splitting/wavelength conversion ability at v is full, we may simply set w(v) = 0. For a node v with 0 < f(v) < 1, its light splitting/wavelength conversion ability is limited. It can be seen that node v has some light splitting/wavelength conversion ability if $0 \le w(v) < 1$. To ensure a multicast tree does not contain any unnecessary node v with w(v) = 0 as its internal node, each node v with w(v) = 0 can be assigned a small value, e.g. $w(v) = \epsilon = \frac{1}{n+1}$, where n is the number of nodes in the network.

2.2 Node Cost Model

For a given WDM network G and a multicast request, a multicast tree rooted at the source and spanning all the terminals is built if there are sufficient network

resources to realize the request. Since each leaf node in the multicast tree only receives messages from its parent, no light splitting/wavelength conversion is needed at the node. Thus, we define the cost C(T) of a multicast tree T in G is the weighted sum of all the internal nodes in T. We refer to this model as the node cost model, which aims to be used in minimizing the utilization of the light splitting/wavelength conversion resources in the multicast tree per request.

2.3 Problem Definition

The multicast tree for a given multicast request (s; D) in G is such a tree rooted at s and spanning all the nodes in D that all its leaf nodes are terminals, where the source $s \in V - D$ and the terminal set $D \subset V$.

The optimal multicast tree for a given multicast request (s; D) is such a multicast tree that the weighted sum of the internal nodes in the tree is minimized.

The optimal multicast tree problem is to find an optimal multicast tree for a given multicast request (s; D) in G. The optimal multicast tree problem is referred to as the optimal broadcast tree problem when $D = V - \{s\}$.

The online multicast request maximization problem for a sequence of multicast requests is to maximize the number of the realized requests in the sequence until the system is unable to accommodate any further requests.

Due to the nature of unforeseen future requests, it is very difficult to provide an exact solution to the online multicast request maximization problem, instead, in this paper we focus on finding a "cost-optimal" multicast tree for each request under the node cost model. We must mention that we here deal with the WDM networks with shared light splitter bank, the availability of light splitters/wavelength converters at a node is the major concern and the link traffic load will not be taken into account in the node cost modelling.

3 Algorithms Based on the Node Cost Model

In this section we first show that the optimal multicast tree problem under the proposed node cost model is NP-complete. We then provide approximation and heuristic algorithms for the problem of concern.

3.1 NP-Hardness of the Optimal Multicast Tree Problem

In the following we show that the optimal broadcast tree problem is NP-complete by a reduction from the maximum leaf spanning tree problem (MLST for short) in G, which is to find a spanning tree in G such that the number of leaves in the tree is maximized. MLST has been shown to be NP-complete [12]. In fact, in terms of computational hardness, the optimal broadcast tree problem and MLST are equivalent within polynomial time. In addition, the optimal broadcast tree problem is a special case of the optimal multicast tree problem, thus, the optimal multicast tree problem is also NP-complete. **Theorem 1.** The optimal broadcast tree problem in a WDM network G(V, E, w) with shared splitter bank is not only NP-complete but also complete for MAX-SNP.

Proof. Given an instance G(V, E) of MLST and an integer k, the decision version of MLST is to determine whether there is a spanning tree in G such that the number of leaf nodes in the tree is no less than k.

We now construct an instance – a WDM network G(V, E, w) of the optimal broadcast tree problem, where each node v in V has identical light splitting/ wavelength conversion ability w(v) = r > 0. Let T_{opt} be the optimal broadcast tree in G and n_1 the number of leaf nodes in T_{opt} . Then, the weighted sum of the internal nodes in the tree is $r * (n - n_1)$, which is the minimum when n_1 is maximized.

Given the instance G(V, E) of MLST, we can see that there is a corresponding instance G(V, E, w) of the optimal broadcast tree problem with an integer r * (n-k), and thus there is a broadcast tree in G(V, E, w) such that the weighted sum of its internal nodes is no more than r * (n-k).

Clearly, to verify whether a given tree is a solution to the optimal broadcast tree problem can be done within polynomial time. Thus, the optimal broadcast tree problem is NP-complete. It is easy to show that the optimal broadcast tree problem and the MLST problem are equivalent in terms of computational complexity under polynomial time reduction. It has been shown that the MLST problem is not only NP-complete [12] but also complete for MAX-SNP [13], which means that it does not permit a fully polynomial-time approximation schema unless P = NP [14]. Thus, it is unlikely to have a fully approximation schema to the optimal broadcast tree problem unless P = NP.

3.2 A Simple Approximation Algorithm

Due to the NP hardness of the optimal multicast tree problem, we provide a simple approximation algorithm for the optimal multicast tree problem, which is referred to as *algorithm* SA.

The edge-weighted directed Steiner tree problem for a source s and a terminal set D is to find a tree in G rooted at s and spanning all the nodes in D such that the weighted sum of the edges in the tree is minimized.

Now we approach the optimal multicast tree problem by reducing it to the edge-weighted directed Steiner tree problem for the source s' and the terminal set D' in an auxiliary directed graph G'(V', E', w'), where $V' = \{v_1, v_2 \mid v \in V\}$, $E' = \{\langle v_1, v_2 \rangle \mid v \in V\} \cup \{\langle v_2, u_1 \rangle, \langle u_2, v_1 \rangle \mid (u, v) \in E\}, w'(\langle v_1, v_2 \rangle) = w(v)$ and $w'(\langle v_2, u_1 \rangle) = w'(\langle u_2, v_1 \rangle) = 0, s' = s_1, D' = \{v_1 \mid v \in D\}.$

Theorem 2. Given a WDM network G(V, E, w), a source s and a terminal set $D, s \in V - D, D \subset V$, assume that G'(V', E', w') is the corresponding auxiliary graph of G. Let T' be a solution to the edge-weighted directed Steiner tree problem for s' and D' in G'. Then, T is a solution to the optimal multicast tree problem for (s; D), where $V(T) = \{v \mid v_1 \in V(T')\}$ and $E(T) = \{\langle v, u \rangle \mid v_1 \text{ is an internal node in } T' \text{ and } \langle v_2, u_1 \rangle \in E(T')\}.$ *Proof.* If v_1 is an internal node in T', then $v_2 \in V(T')$ because $\langle v_1, v_2 \rangle$ is the only edge starting from v_1 in G'. Since $v_2 \notin D'$, v_2 is not a leaf node in T', then there exists a node u_1 in T' such that $\langle v_2, u_1 \rangle \in E(T')$. Thus, T is a tree.

If $\langle v_1, v_2 \rangle$ is an edge in T', v is an internal node in T. Then we have C(T) = W(T'), where C(T) is the weighted sum of the internal nodes in T, and W(T') is the weighted sum of the edges in T'. Now we prove that C(T) is minimized. If there is another tree T_1 rooted at s and spanning all the nodes in D and $C(T_1) < C(T)$. We define T'_1 as follows.

$$V(T'_1) = \{v_1 \mid v \in V(T_1)\} \cup \{v_2 \mid v \text{ is an internal node in } T_1\}$$
$$E(T'_1) = \{\langle v_1, v_2 \rangle, \langle v_2, u_1 \rangle \mid \langle v, u \rangle \in E(T_1)\}$$

Then, T'_1 is a tree in G' rooted at s' and spanning all the nodes in D', and $W(T'_1) = C(T_1)$. Thus, we have $W(T'_1) = C(T_1) < C(T) = W(T')$, which contradicts to the assumption that T' is a solution to the edge-weighted directed Steiner tree problem for s' and D'.

Following Theorem 2, an approximation solution to the edge-weighted directed Steiner tree problem in G' can be transformed into an approximation solution to the optimal multicast tree problem in G. It is known that the best possible approximation solution for the directed Steiner tree problem so far is $O(|D'|^{\delta})$ times of the optimum [15], where δ is a constant with $0 < \delta \leq 1$, |D'| = |D|. We thus have the following theorem.

Theorem 3. Given a WDM network G(V, E, w) with shared light splitter bank and a multicast request (s; D), there is an approximation solution to the optimal multicast tree problem, which is $O(|D|^{\delta})$ times of the optimum, where δ is a constant with $0 < \delta \leq 1$.

3.3 A Heuristic Algorithm

In the following we propose a heuristic algorithm for the optimal multicast tree problem. The proposed heuristic is similar to the approximation algorithm for the node-weighted Steiner tree problem, referred to as *algorithm* KR, by Klein and Ravi [16] but with some important modifications. The node-weighted Steiner tree problem is to find a tree in G spanning all the nodes in terminal set D such that the weighted sum of the nodes in the tree is minimized.

The algorithm KR maintains a forest F that consists of a node-disjoint set $\{T_1, T_2, \ldots, T_k\}$ of trees and contains all the terminals, $1 \le k \le |D|$. Initially, each terminal by itself is a tree. The algorithm uses a greedy strategy to iteratively merge several current trees into a larger tree until there is only one tree left in the forest. In each iteration, the algorithm selects a node and a subset of the current trees of size at least two so as to minimize the ratio

$$\frac{w(v) + \sum_{T_j \in S} d(v, T_j)}{|S|} \tag{1}$$

where $S \subseteq F, |S| \ge 2$, $d(v, T_j)$ is the distance from v to T_j . The distance along a path in algorithm KR does not include the weights of the two endpoints of the path. Thus, the choice minimizes the average node-to-tree distance. The algorithm uses the shortest paths between the node and the selected trees to merge the trees into a larger one.

To implement each iteration, for each node v, the quotient cost of v is defined to be the minimum value of (1), taken over all subsets of the current trees of size at least two. To find the quotient cost of v, the algorithm computes the distance d_j from v to each T_j , assuming without loss of the generality that the trees are numbered so that $d_1 \leq d_2 \leq \ldots \leq d_k$. In computing the quotient cost of v, it is sufficient to consider subsets of the form $\{T_1, T_2, \ldots, T_i\}, 2 \leq i \leq k$. The quotient cost for a given node can be calculated in polynomial time, and the minimum quotient cost can then be determined. Thus, each iteration can be carried out within polynomial time. The solution delivered by the algorithm for node-weighted Steiner tree problem is $2 \ln |D|$ times of the optimum. Note that the approximation of the solution is within a constant factor of the best possible approximation achievable in polynomial time unless $\tilde{P} \supseteq NP$ [17]. Guha and Khuller [18] later provided an improved algorithm for the problem with a better approximation ratio at the expense of a longer running time. Their improved algorithm delivers a solution within 1.35 $\ln |D|$ times of the optimum.

It should be emphasized that the problem we deal with is different from the one discussed by Klein and Ravi [16], despite there being some similarities between them. Their approximation analysis is based on an assumption that the weight of each terminal is zero, since all the terminals will be included into the Steiner tree. However, for the optimal multicast tree problem, we treat each terminal differently, depending on whether or not it is an internal node in a multicast tree. If it is, its node weight should be taken into account; otherwise its node weight can be ignored because a leaf node is only a receiver of messages and no light splitting/wavelength conversion is needed at it. Thus, the solution delivered by algorithm KR is not an approximation solution for the optimal multicast tree problem. Now, we propose a heuristic for the optimal multicast tree problem based on some modifications to algorithm KR. The differences between our heuristic and algorithm KR are at the following crucial steps in defining the length of a path between two nodes and calculating the quotient cost of a node.

Assume that there are k trees T_1, T_2, \ldots, T_k currently, $k \leq |D|$. To compute the quotient cost of a given node v, we need to compute the distance from vto T_j , which is in turn reduced to computing the *length* of the shortest paths between v and every node u in T_j . In our algorithm, the length of a path between v and a tree node u is the weighted sum of the nodes in the path except u and v if u is not a leaf node in T_j or $T_j = \{u\}$; otherwise, the length of the path is the weighted sum of all the nodes in the path except v. While computing the quotient cost of a node v in (1) in our algorithm, w(v) is not taken into account if v is an internal node in one of the current trees; otherwise w(v) is included in the calculation of the quotient cost of v. We refer to this heuristic as *algorithm* MKR.

4 Performance Study

In this section we evaluate the performance of the proposed approximation algorithm SA and heuristic algorithm MKR against that of two existing algorithms KR and SPT in terms of network throughput by conducting experimental simulations, where SPT is the edge-weighted shortest path tree algorithm in which each edge has identical weight.

4.1 Simulation Environment

We assume that 100 nodes are deployed randomly in a region of $10 \times 10 \text{ m}^2$ using the NS-2 simulator. For each pair of nodes u and v, a random number $r_{u,v}$ is generated, $0 \leq r_{u,v} < 1$. Whether or not u and v are connected is determined by $r_{u,v}$ and the edge probability [19,20]

$$P(u,v) = \beta e^{\frac{-d(u,v)}{L\alpha}}$$

where d(u, v) is the Euclidean distance between u and v, L is the maximum distance between nodes in the region, and α and β are the parameters governing the edge density in the network, $0 < \alpha, \beta \leq 1$. There is an edge between u and v if and only if $r_{u,v} < P(u, v)$. Different values of α and β result in different network topologies even with the same node distribution.

We assign a weight to each node in the network to model its light splitting/wavelength conversion ability. Initially, the weight w(v) is a random number between zero and one. w(v) increases by c if v is an internal node of the multicast tree built for a multicast request that consumes light splitting/wavelength



Fig. 1. Comparison of the network throughputs with $\alpha = 0.3$ and $\beta = 0.3$



Fig. 2. Comparison of the network throughputs with various values of α and β

conversion resources of amount c. Node v has no light splitting/wavelength conversion ability when its current weight is greater than or equal to one.

We assume that the sequence of multicast requests consists of 200 requests. We randomly select the source and the terminal set for a multicast request, and the size of the terminal set ranges from 10 to 50 with increments of 10. We also assume that each multicast request lasts for a period of time and consumes a certain amount of light splitting/wavelength conversion resources. For simplicity, we further assume that the consumption of these resources is identical for all the internal nodes in a multicast tree.

We simulated various algorithms on 10 different randomly generated network topologies for different problem size. For each size of the network instance, the value shown in the graphs is the mean of 10 individual values obtained by running each algorithm on these 10 randomly generated network topologies.

4.2 Performance Evaluation of Various Algorithms

We first evaluate the network throughputs delivered by different algorithms with various sizes of the terminal sets when $\alpha = 0.3$ and $\beta = 0.3$. It can be seen from Fig. 1 that algorithms MKR and SA outperform algorithms KR and SPT significantly with the sizes of the terminal sets varying from 10 to 50 with increments of 10. When the terminal set consists of 10 nodes, more than 90% of the requests can be realized by algorithms MKR and SA, whereas the realization ratios of algorithms KR and SPT are only 80% and 60% respectively. In addition, the network throughput delivered by algorithm KR drops faster than those delivered by the other algorithms MKR and SA can still realize around 70% of the requests. It also can be observed from Fig. 1 that the network throughput delivered by algorithm KR throughput delivered by algorithm SA for various sizes of terminal sets.

We then change the edge density in various network topologies by varying the values of α and β . It is indicated in Fig. 2 that there is no significant difference among the algorithms in terms of the performance, compared with the case where $\alpha = 0.3$ and $\beta = 0.3$, i.e. the performance of algorithms MKR and SA is constantly better than that of algorithms KR and SPT.

5 Conclusion

In this paper we have studied online multicasting in WDM networks with shared light splitter bank aiming at maximizing the network throughput. We first proposed a node cost model for multicast trees that models the cost of utilizing the network resources such as light splitters/wavelength converters at nodes. We then showed that finding a cost-optimal multicast tree under the proposed cost model is NP-complete, and instead devised approximation and heuristic algorithms for finding such cost-optimal multicast trees. We finally conducted experiments to evaluate the performance of the proposed algorithms. The experimental results show that the proposed algorithms are efficient and effective in terms of network throughput.

Acknowledgment. It is acknowledged that the work by the authors is fully funded by a research grant No:DP0449431 by Australian Research Council under its Discovery Schemes.

References

- T. F. Znati, T. Alrabiah and R. Melhem. Low-cost delay-bounded point-to-point multipoint communication to support multicasting over WDM networks. *Computer Networks*, Vol.38, pp.423–445, 2002.
- G. N. Rouskas. Optical layer multicast: rational, building blocks, and challenges. IEEE Network, Vol.17, pp.60–65, 2003.
- M. Ali and J. Deogun. Power-efficient design of multicast wavelength-routed networks. *IEEE JSAC*, Vol.18, pp.1852–1862, 2000.
- 4. Y. Zhou and G. S. Poo. Optical multicast over wavelength-routed WDM networks: A survey. *Optical Switching and Networking*, Vol.2, pp.176–197, 2005.
- Z. Zhang and Y. Yang. Online optimal wavelength assignment in WDM networks with shared wavelength converter pool. *Proc. of IEEE INFOCOM'05*, pp.694–705, 2005.
- G. Sahin and M. Azizoglu. Multicast routing and wavelength assignment in widearea networks. *Proc. of SPIE Conf. on All-Optical Networks*, Vol.3531, pp.196–208, 1998.
- R. Libeskind-Hadas and R. Melhem. Multicast routing and wavelength assignment in multi-hop optical networks. *IEEE/ACM Trans. Networking*, Vol.10, pp.621–629, 2002.
- J -C Bermond, L. Gargano, S. Perennes, A. Rescigno, and U. Vaccaro. Efficient collective communications in optical networks. *Theoretical Computer Science*, Vol.233, pp.165–189, 2000.
- 9. R. Libeskind-Hadas. Efficient collective communication in WDM networks. *Proc.* of IC^3N , pp.612–616, 2000.
- 10. X. Zhang, J. Wei and C. Qiao. Constrained multicast routing in WDM networks with sparse light splitting. *Proc. of IEEE INFOCOM'00*, pp.1781–1790, 2000.
- B. Chen and J. Wang. Efficient routing and wavelength assignment for multicast in WDM networks. *IEEE Journal on Selected Areas in Communications*, Vol.20, pp.97–109, 2002.
- M. R. Garey and D. S. Johnson. Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman, 1979.
- G. Galbiati, F. Maffioli and A. Morzenti. A short note on the approximability of the maximum leaves spanning tree problem. *Inform. Proc. Lett.*, Vol.52, pp.45–49, 1994.
- C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. Proc. of STOC, ACM, pp.229–234, 1988.
- M. Charikar, C. Chekuri, T-Y Cheung, Z. Dai, A. Goel, S. Guha and M. Li. Approximation algorithms for directed Steiner problems. J. Algorithms, Vol.33, pp.73–91, 1999.

- P. N. Klein and R. Ravi. A nearly best-possible approximation algorithm for nodeweighted Steiner trees. J. Algorithms, Vol.19, pp.104–114, 1995.
- C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. J. of the ACM, Vol.41, pp.960–981, 1994.
- S. Guha and S. Khuller. Improved methods for approximating node weighted Steiner trees and connected dominating sets. *Inform. and Comp.*, Vol.150, pp.57– 74, 1999.
- B. Waxman. Routing of multipoint connections. *IEEE Journal on Selected Areas in Communications*, Vol.6, pp.1617–1622, 1988.
- F. Bauer and A. Varma. ARIES: a rearrangeable inexpensive edge-based online Steiner algorithm. *IEEE Journal on Selected Areas in Communications*, Vol.15, pp.382–397, 1997.