Minimum Interference Optical Multicast

FRANCESCO PALMIERI and UGO FIORE
Centro Servizi Didattico Scientifico
Università degli studi di Napoli "Federico II"
Complesso Universitario di Monte S. Angelo,Via Cinthia 5
NAPOLI - ITALY

Abstract: - All-optical networks with wavelength division multiplexing and wavelength-routing offer an excellent application environment for multicast applications such as video-conferencing, high-definition TV (HTDV), distributed simulation and content distribution in mirror sites and in Storage Area Networks, that can operate entirely in the optical domain. Accordingly, native multicast support in next generation all-optical networks is an extremely challenging topic. In this paper we focus on the efficient construction of light-forests under optical-layer power budget and wavelength continuity constraints. We introduced a novel heuristic to improve network utilization in multicast RWA, aiming at addressing the problem of reducing blocking probability for multicast group when building multiple light-trees in a light-forest. In brief, we extended to the multicast environment the idea of minimum interference, well known in bandwidth routing research and avoided any links if they minimize the maximum flow available for future groups. We also studied the performance of the proposed heuristic by conducting simulation experiments which demonstrated its efficiency.

Key-Words: - Optical networks, multicast, maximum flow, routing and wavelength assignment.

1 Introduction

As the internet traffic demand continues to increase exponentially, wavelength division multiplexing (WDM), providing terabits per second bandwidth per fiber by multiplexing many different wavelength signals onto a single optical strand, becomes the natural choice for the backbone technology in the next generation Internet. The wavelength routing technology is considered extremely promising for the realization of networks that will have to handle a huge IP data traffic, including both unicast and multicast. Particularly, since IP multicast is increasingly popular and strategic on the Internet, because it is the only efficient and scalable way to provide one-to-many or many-to-many communications, issues concerning its support over next-generation wavelength routed networks become an important and yet challenging topic.

In such networks routing unicast connections requires determining a path, often referred as lightpath, from the source of the optical signal to its destination, which may span a number of fiber links. In addition to determining the route for a connection, a wavelength must be assigned to that connection such that all other connections which share some fiber link with this route are assigned different wavelengths. If no wavelength converters are used, a lightpath will be realized with the same wavelength on each hop, according to the wavelength continuity constraint. On the other side, when using converters, that are known to be expensive and often poor in performance and error prone, a lightpath may be created by using a different wavelength on each hop where such a device is available.

Similarly, whilst optical unicast connections carried on lightpaths are point-to-point connections and since routing a multicast connection requires the discovery of a tree rooted at the source which spans all destinations, multicast in the optical domain, requires point-to-multipoint non looping wavelength connections which may only be described by trees, rooted at the source node, often referred as light-trees, that generalize the concept of lightpath [1]. The light-tree, like a lightpath, is a clear channel originating at a given source node and implemented, in presence of the wavelength continuity constraint, with a single wavelength, but, unlike a lightpath, it may have, as a point-to-multipoint channel, multiple destination nodes.

Native multicast is supported at the optical layer by letting WDM switches or OXCs make copies of data packets in the optical domain via light splitting. In detail, to provide point-to-multipoint connections, each intermediate node that has more than one child in a light-tree must be equipped with an optical power splitter dividing the power of the input signal into more than one output signal; thus reducing the input signals power by a factor proportional to the number of outputs. This scheme is desirable since only one laser is needed for transmission and links can be shared in the tree structure. On the other hand, an optical signal must have a minimum amount of power in order to be dropped-off at a destination or passed to the next downstream node. Therefore, due to this split loss, it is not possible to drop off data at an arbitrary number of destinations in a single light-tree. This is commonly known as the optical power budget constraint. Clearly, each optical network node must have, to be multicast capable, enough light splitting capability, and lightpaths have to be split if needed, even when the paths to different nodes do not diverge, i.e. it is possible to have multiple multicast tree branches on the same link.

Although the problem of establishing a light-tree that spans a given source and a set of destination nodes bears some similarities to the well known Steiner tree problem, the nature of optical multicast introduces several new issues and complexities. Splitting an optical signal introduces losses, a problem not encountered in electronic packet- or circuit-switched networks, and thus, not addressed by existing routing tree algorithms. Even in the presence of optical amplifiers, this signal loss imposes a hard upper bound on the number of times a signal can be split, as well as on the number of hops that the signal can travel after every split operation.

In the absence of wavelength conversion in the network (or even in networks with limited or sparse conversion capability), multicast routing is tightly coupled to wavelength allocation, an issue that does not arise in electronic networks. Also, optical networks may have only a sparse multicast switching capability; i.e., only a subset of the optical switching devices may be multicast capable. When only a few multicast capable switching nodes are present in the network, a feasible multicast tree may not exist, and therefore the heuristics for degree-constrained multicast are not applicable at all.

Finally, the capacity planning problems related to placing multicast capable optical switching devices in the network [2] and multicast routing strongly depend on each other. Several recent research efforts have aimed at addressing some of the problems associated with optical multicast and light-tree establishment, including studies of wavelength assignment in the presence of multicast [3][4] and multicast routing algorithms for networks with a sparse light splitting capability [5][6]. To deal with the fact that a feasible multicast tree may not exist for a given source and destination set, the concept of a light-forest has been proposed [5]. In general, all the multicast routing algorithms for optical networks assume unlimited fanout capacity at multicast capable nodes, and each tree of a given light-forest must be assigned a different wavelength.

In this paper we focused on the efficient construction of light-forests under optical-layer power budget and wavelength continuity constraints. We investigated the idea of using the minimum maximum flow between some subset of the possible source nodes and destinations nodes as the quantity to maximize, in order to get a routing and wavelength assignment paradigm that preserves as much connectivity as possible. This will, in turn, reduce the blocking probability of future requests. The idea, initially proposed by Kodialam and Lakshman with the Minimum Interference Routing Algorithm, is well known in the context of bandwidth routing, and we transposed it into the context of optical multicast by introducing a novel heuristic that seems to be effective in improving network utilization and reducing blocking probability for multicast groups. We studied the performance of the proposed heuristic by conducting simulation experiments which demonstrated its efficiency.

2 Related work

There has been much study focusing on the optimal RWA for a multicasting request, possibly under various constraints such as sparse light splitting and limited light power. Libeskind-Hadas [7] used a *k-drop multi-tree model* to take power budget constraints into account. Under that model, at most a pre-specified number *k* of destination nodes can receive data in each light-tree, so the problem becomes that of finding a minimal-cost *set* of light-trees spanning the set of destination nodes. Wavelength assignment must guarantee that distinct wavelengths are used for two light-trees sharing a common link. In [8], a greedy algorithm is used, to find a light-forest, such that each light-tree uses the same wavelength and the total cost in minimized. Zhang *et al.* proposed in [5] four algorithms to construct a source-based multicast light-forest consisting of one or more multicast trees. The objectives of the algorithms include minimizing the number of wavelengths and the number of hops from the source to each destination. He *et al.* [9] classify multicast groups by assigning weights to them, and use that extra information in a RWA heuristic that aims to minimize the weighted overall blocking rate.

Many multicast tree formation algorithms, which construct a source-based tree given the full knowledge of network topology and multicast session membership, have been proposed and their performance evaluated in the literature (see, e.g., [10] for a survey). These heuristic algorithms can roughly be classified into three categories. The first one contains algorithms based on the shortest path heuristic (SPH) which initializes the Steiner tree to the shortest path from the source to an arbitrary multicast member the cost of the path from a multicast source to each of the members. Then, new members are repeatedly added by selecting the shortest path between the member to the current partial tree, until all members have joined the tree. The effect on the final tree quality of ordering strategies that drive the member inclusion process, such as including the members in the order determined by their distance to the multicast tree instead of random inclusion, or growing the Steiner tree from the destinations instead of from the source, has also been investigated.

The second one contains algorithms based on the minimum Steiner tree, which attempt to minimize the total cost of a multicast tree, by using a graph with as many vertices as they are involved in the multicast group. This is achieved by constructing a closure graph of the multicast nodes from the original graph, using the cost of the shortest path between each pair of members. Polynomial-time algorithms give the minimum spanning tree of the closure graph. To obtain the multicast tree, the shortest paths in the original graph are used to replace the edges of the minimum spanning tree, and cycles are removed. The last category embraces simulated annealing, genetic algorithms, and Tabu search, that have been investigated to solve the Steiner tree problem and have been shown to perform well on average.

3 Problem statement

The generic multicast delivery problem in communication networks is described by the Steiner Tree Problem, defined as follows. Given a graph *G(V,E)* where *V* is the set of vertices and *E* the set of edges; a cost function $c : E \rightarrow \Re^+$ applied to each edge of *E*; and a set of nodes $N \subset V$; find a sub-tree *T = (V[T], E[T])* spanning *N* and such that its cost *c(T)* (defined as the sum of individual edge cost *c(e)* for all the edges $e \in E[T]$) is minimized. The Steiner Tree problem, also if not restricted to the optical domain, is NP-Complete [1] and hence, even if a method to optimize the multicast tree topology of an optical network exists in theory, in practice it is not possible to analytically solve this optimization problem and a heuristic approach is required to setup multicast trees in an efficient way.

For instance, in packet-based multicast networks, shortest-path trees are computed in order to implement a fully distributed on-line computation of the multicast-tree. On the other side, before a multicast packet belonging to a specific traffic flow has to be forwarded in an all-optical wavelength routed network it has to be assigned a wavelength and hence routed on a lightpath. Since there are many wavelengths on each outgoing fiber, scheduling has to be performed to find out which wavelength is available at what time. If no wavelength is available at the time the packet should be forwarded, it has to be either buffered for transmission at a later time (maybe on a different wavelength) or dropped.

The problem of optimally assigning wavelengths and routes for a set of connections is referred to as Routing and Wavelength Assignment (RWA) and is known to be NP-complete even in trivial network topologies. That's worse, in WDM networks, optical buffer space is very limited or absent at all, while current IP multicast routing protocols are designed under the assumption that buffer space is unlimited (or very large). Buffer space limitation has a significant impact on the blocking/dropping probability. For multicast traffic, the buffer space limitation may also affect multicast routing. Partial failure will occur in case of congestion on any subset of downstream interfaces. How to deal with congestion (e.g. by rerouting around some congested branches or by using intelligent multicast routing algorithm to balance the traffic load on each link) is a challenging issue. Hereafter, we will focus on the routing problem when supporting IP multicast at the WDM layer assuming that only a subset of switches is multicast capable.

If the traffic patterns in the network are reasonably well known in advance and any traffic variations take place over long timescales, the most effective technique for establishing optical connections (lightpaths or light-trees) is by solving a static multicast RWA (MC-RWA) problem. Since these connections are assumed to remain in place for relatively long periods of time, it is worthwhile to optimize the way in which network resources (e.g., physical links and wavelengths) are assigned to each connection, even though optimization may require considerable computational effort.

Typically, routing and wavelength assignment are considered together as an optimization problem using integer programming formulations. An instance of the static MC-RWA problem with the objective of minimizing the number of wavelengths was considered in [3], and a set of heuristics was presented. The dynamic MC-RWA problem is encountered during the real-time network operation phase, and involves the dynamic provisioning of light-trees or light-forests. Specifically, users submit to the network, in some random fashion, requests for light-trees to be set up as needed. Depending on the state of the network at the time of the request, the available resources may or may not be sufficient to establish a light-tree that spans the requested source node and destination set. The network state consists of the physical route and wavelength assignment for all active light-trees, and it evolves randomly in time as new light-trees are admitted and existing ones released. Thus, each time a request is made, the network must determine whether it is feasible to accommodate the request, and if so, to perform routing and wavelength assignment. If a request for a light-tree cannot be accepted due to lack of resources, it is blocked. Because of the real-time nature of the problem, MC-RWA algorithms in a dynamic traffic environment must be simple and fast.

Since combined routing and wavelength assignment is a hard problem, a typical approach to designing efficient algorithms is to decouple the problem into two separate subproblems: the light-tree routing problem and the wavelength assignment problem. This approach has been taken in [4][11][12] to study the blocking performance of optical networks with dynamic multicast traffic. All three studies decouple the routing and wavelength assignment problems, and consider alternate routing strategies, whereby a number of trees may be used for each multicast connection. When a request for a connection arrives, the associated trees are considered in some order, and the connection is blocked if no free wavelength is found in any of the trees. The study in [4] considered the first-fit wavelength allocation policy, in which for each tree, wavelengths are also considered in a fixed order. In [12], on the other hand, the random wavelength assignment was studied; in addition, multiple classes of requests were considered. A different approach was taken in [11], where an iterative approximation algorithm was developed for completely connected networks under random wavelength assignment. Since, in general, network topologies are not completely connected, the results of [11] can be used as lower bounds for more general topologies. All these inherent complexities and constraints make the construction of efficient WDM multicast trees somewhat difficult and may possibly lead to trees that are different from their corresponding ones in the IP layer.

4 System model

The WDM network is modeled as an undirected graph $G = (V,E)$, where V is the set of nodes with $|V| = N$ and E is the set of links. The nodes are labeled $v_1, v_2, \ldots, v_N$. All of the nodes have light-splitting capability. The set of wavelengths available on each link is $L = \{\lambda_1, \lambda_2, \ldots, \lambda_W\}$, where W is the total number of wavelengths. There is no wavelength conversion capability at the network nodes, and hence the wavelength continuity must be met.

There are M multicast groups. Each multicast group is represented as $g_i = (s_i, D_i)$, where $s_i \in V$ is the source node and $D_i \subseteq V - \{s_i\}$ is the set of destination nodes of the multicast group i. In the scenario we are considering, group requests arrive one at a time.

The RWA problem is to determine the route and wavelength assignment for each group such that the overall blocking rate can be minimized given the group distributions $g_i = (s_i, D_i)$, $i = 1, \ldots, M$. The light-forest for g_i, $F_i(s_i, D_i)$, consists of t_i multicast trees rooted at s_i, $T_{ij}(s_i, D_{ij})$, with $\bigcup_{j=1}^{t_i} D_{ij} = D_i$ and

$D_{ij} \cap D_{ik} = \varnothing$ if $k \neq j$. We denote the RWA for group i as a collection of tuples $\left(T_{ij}, \lambda'_{ij}\right)$, where $T_{ij} \subseteq E$ is the set of links used to serve the j-th tree in the forest and $\lambda'_{ij} \in L$ is the wavelength used on these links (all these links use the same wavelength due to wavelength continuity).

Because the multicast RWA problem is NP-complete [3], we propose a heuristic to minimize the overall blocking rate in this section. Note that the aforementioned RWA problem consists of two basic subproblems, namely, the routing problem (determining T_{ij}) and the wavelength assignment problem (determining λ'_{ij}). To solve the routing problem, a multicast tree is built on each wavelength graph (a subgraph of G obtained by removing the links on which the wavelength is not available) and the best one is chosen. After a group is served, the corresponding wavelength graph is reduced due to the removal of the used links. In this process some nodes would be disconnected in the wavelength graph because all the links incident on them have been removed. If this particular wavelength is used to serve another group later, only a subset of the destination nodes can be served. In other words, the multicast groups served later are more likely to be blocked.

5 The MaxFlow Heuristic

We investigate in this study the idea of using the minimum maximum flow between some subset of the possible source nodes and destinations nodes as the quantity to maximize, in order to get an RWA that preserves as much connectivity as possible. This will, in turn, reduce the blocking probability of future requests. The idea is well-known in the context of bandwidth routing. It was initially proposed by Kodialam and Lakshman [13], and was further developed in by Suri *et al.* [14]. We are unaware, though, of any attempts to apply it to optical multicasting in the presence of limited splitting capacity. The basic observation is that serving a multicast group on a light-tree can reduce the available capacity for other groups. The maximum permissible flow to other groups is a measure of that capacity. If trees that decrease the possible maxflow between other nodes by a large amount are avoided, the creation of some bottlenecks can also be avoided. Instead of being single source-single sink, the maxflow problems to be addressed in the case of optical multicast are single source-multiple sink.

Function MaxFlow
BEGIN
  Select virtual groups Y
  $\forall g_i \in Y$ **DO**
    Compute Max Flow for g_i
  END
  Adjust individual link weights,
  according to the flows computed before.

  Compute minimum-cost light-forest
END

Figure 1. The MaxFlow heuristic.

Actually solving the maxflow problems for every possible multicast group is not an option, because the number of possible receiver subsets is exponential in the size of V. To reduce the time complexity of the computation, we probabilistically choose some multicast groups. While an *a priori* knowledge of the entire multicast traffic is unrealistic, the traffic patterns can be estimated by using previous patterns as well as forecasts based on additional knowledge of the network usage. A virtual multicast traffic scheme, representing a realistic sample of future requests is thus computed, by using statistical information about multicast traffic offered. After forecasting upcoming multicast groups in the immediate future, those groups that are likely to interfere with the current group are selected.

6 Simulation and results

To facilitate performance comparison among the proposed light-tree construction algorithms, we assume each link has a sufficient number of wavelengths to avoid blocking. In addition, the First-Fit algorithm will be used (although other heuristics may also be used) to perform wavelength assignment after the light-trees are constructed and partitioned into segments. We will determine the maximum number of wavelengths needed by a given forest (over all the links), and then use the average maximum number of wavelengths needed per forest (over many forests and simulation runs), denoted by W, which represents the amount of network resources required per forest, as the first performance metric. In wavelength-routed WDM networks, one wavelength channel (or a unit of bandwidth) needs to be reserved on each branch of a light-forest. For simplicity, we assume that all wavelengths are equally expensive (or cheap), and in addition, the bandwidth consumed using a wavelength on different links is more or less the same as well.

Accordingly, we will determine the (total) number of branches on a multicast forest, and then use the average number of branches per forest, denoted by *B*, which represents the average bandwidth consumed per forest, as the second performance metric. Finally, for a given forest (and multicast session), we will determine the average number of hops from the multicast source to a destination (over all the destinations of the multicast session), and then use the averaged value over many forests and simulation runs, denoted by *H*, which represents the delay, as the third performance metric.

6.1 Simulation

In the simulations the backbone of NSFnet [12] was chosen as the network topology. It consists of 16 nodes representing states in the USA (Figure 2). The network cost of a link joining two states is the driving distance between them.

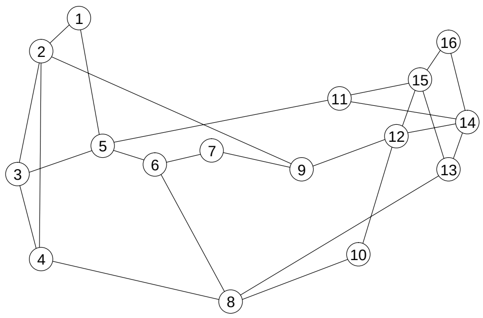

Figure 2. Simulation topology.

The multicast groups were randomly generated, with two separate sets of constraints. We expected the MaxFlow Heuristic it to perform poorly in the presence of widely different random multicast groups. However, we also simulated random multicast connections, to evaluate the performance of the strategy under very difficult conditions. The first batch of multicast connections are generated according to the following hypotheses: four of nodes act as sources; groups span from five to all the fourteen nodes. Groups in the second pool are obtained by randomly picking up from the nodes in the networks the source node s and set D of destination nodes.

The performance of the MaxFlow heuristic was compared against the SPT heuristic.

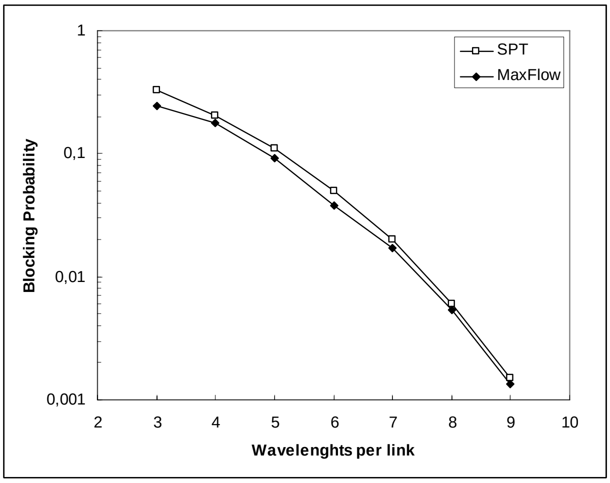

Figure 3 Blocking probability – constrained groups

In the simulations, the number of blocked requests were simulated against three parameters: number of wavelengths, average number of multicast destinations, and average number of multicast requests. Due to space restrictions, only the first plots are included.

The blocking probabilities are plotted in Figure 3 as a function of the number of wavelengths per link. MaxFlow outperforms SPT by 20% on the average.

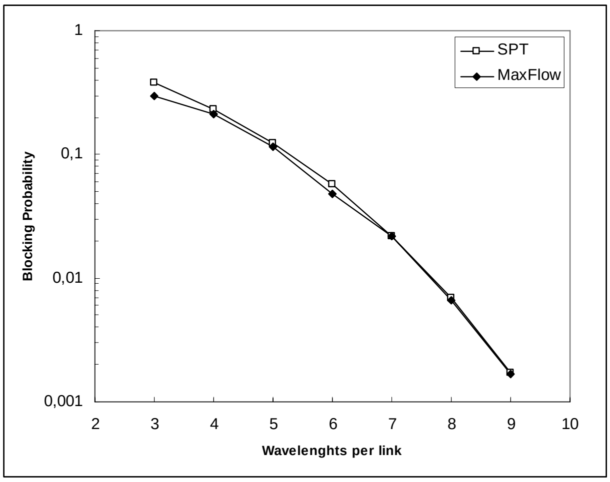

Figure 4. Blocking probabilities – random groups.

Figure 4 shows that MaxFlow performs better than SPT, when the groups are completely random, though the savings are less significant.

7 Conclusion and future work

In this paper, we addressed the problem of reducing the blocking probability for multicast groups in a WDM network. We extended to the multicast environment the idea of minimum interference, well known in bandwidth routing research. In short, links are avoided if they minimize the maximum flow available for future groups. We have studied the performance of the proposed heuristic by conducting simulation experiments. We will work on the improvement of the computational effort required, by isolating individual links that more severely affect the available flow between candidate groups.

References

[1] Sahasrabuddhe, L. H., and Mukherjee, B., "Light-trees: optical multicasting for improved performance in wavelength-routed networks," *IEEE Communications Magazine*, Vol. 37, No. 2, 1999, pp. 67-73.

[2] Ali, M., "Optimization of Splitting Node Placement in Wavelength-Routed Optical Networks," *IEEE Journal on Selected Areas in Communications*, Vol. 20, No. 8, 2002, pp. 1571–1579.

[3] Jia, X.-H., Du, D.-Z., Hu, X.-D., Lee, M.-K., and Gu, J., "Optimization of wavelength assignment for QoS multicast in WDM networks," *IEEE Transactions on Communications*, Vol. 49, 2001, pp. 341-350.

[4] Sahin G., and Azizoglu, M, "Multicast Routing and Wavelength Assignment in Wide-Area Networks," *Proc. SPIE*, vol. 3531, 1998, pp. 196–208.

[5] Zhang, X., Wei, J., and Qiao, C., "Constrained multicast routing in WDM networks with sparse light splitting," in *INFOCOM 2000: Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies*, 2000, pp. 1781–1790.

[6] Tseng; W.-Y., Kuo, S.-Y., "All-optical multicasting on wavelength-routed WDM networks with partial replication," in *Proceedings. 15th International Conference on Information Networking,* 2001, pp. 813-818.

[7] Libeskind-Hadas, R., "Efficient collective communication in WDM networks with a power budget", in *Proceedings of IEEE International Conference on Computer Communication Network*, 2000, pp. 612-616.

[8] Xue, G., and Banka, R., "Bottom-up construction of dynamic multicast trees inWDM networks," in *Proceedings of the IEEE 2003 International Performance, Computing, and Communications Conference*, 2003, pp. 49–56.

[9] He, J, Tsang, H. K., Chan, S.-H., "Multicasting in WDM networks with heterogeneous group weights", *Journal of Optical Networking*, Vol. 3, No. 7, 2004, pp. 534-547.

[10] Rouskas, G. N. "Optical layer multicast: rationale, building blocks, and challenges," *IEEE Netw.*, Vol. 17, No. 1, 2003, pp. 60–65.

[11] Kamal, A. E., and Al-Yatama, A. K., "Blocking Probabilities in Circuit-switched Wavelength Division Multiplexing Networks Under Multicast Service," *Perf. Eval.*, Vol. 47, No. 1, 2002, pp. 43–71.

[12] Ramesh, S., Rouskas, G. N., and Perros, H. G., "Computing Call Blocking Probabilities in Multi-class Wavelength Routing Networks with Multicast Traffic," *IEEE JSAC*, Vol. 20, No. 1, 2002, pp. 89–96.

[13] Kodialam, M., Lakshman, T. V., "Minimum Interference Routing with Applications to MPLS Traffic Engineering," in *Proceedings of INFOCOM 2000*, 2000, pp. 884-893.

[14] Suri, S., Waldvogel, M., Bauer, D., and Warkhede, P. K., "Profile-based routing and traffic engineering," *Computer Communications*, Vol. 24, No. 4, 2003, pp. 351-365.