

Performance Analysis Using A Hierarchical Loss Network Model

Mingyan Liu and John S. Baras

Department of Electrical and Computer Engineering
and Institute for Systems Research
University of Maryland College Park
College Park, MD 20742

Abstract—In this paper we present a hierarchical loss network model for estimating the end-to-end blocking probabilities for large networks. As networks grow in size, nodes tend to form clusters geographically and hierarchical routing schemes are more commonly used. Loss network and reduced load models are often used to approximate end-to-end call blocking probabilities and hence throughput. However so far all work being done in this area is for flat networks with flat routing schemes. We aim at developing a more efficient approximation method for networks that have a natural hierarchy and/or when some form of hierarchical routing policy is used. We present two hierarchical models in detail for fixed hierarchical routing and dynamic hierarchical routing policies, respectively, via the notion of network abstraction, route segmentation, traffic segregation and aggregation. Computation is done separately within each cluster (local) and among clusters (global), and the fixed point is obtained by iteration between local and global computations. We also present numerical results for the first case.

I. INTRODUCTION

Modern networks are getting larger and larger, and it is thus necessary for network modeling and engineering tools to have scale-up capabilities. With the increase in size, a network tends to have clusters of nodes geographically, and hierarchical routing schemes are more commonly used in order to cope with large network size, e.g., ATM PNNI, IP based routing OSPF and hierarchical variations of OSPF. This has provide us with strong motivations as well as applicability to develop a hierarchical model to estimate end-to-end network performances.

In this paper we consider a class of loss network models [1], which is often used to estimate call blocking probabilities, an important performance metric in telephone networks. Various analytical approaches and approximation schemes have been suggested to establish and solve a loss network model, and among them, extensive research has been done in developing and using reduced load/fixed point approximations. With the development of QoS routing and ATM networks, the same technology of reduced load approximation has been applied to packet switched networks for connection level study via the concept of effective bandwidth [2].

While research results are abundant for fully connected, symmetric networks with fixed, sequential or state-dependent routing [1], especially for networks with no more than two hops on their routes [3], or when network traffic is of single rate [4], there has been far less attention to large random networks with

both multiple traffic rates and state-dependent routing [1], [5], [6], [7]. Furthermore, all of such methods are for flat networks and flat routing schemes.

Motivated by the above, we develop a hierarchical version of the reduced load model. The objective is not to artificially introduce hierarchy into a network model, but rather to develop a more efficient and scalable way of performance analysis for networks that have a natural hierarchy and/or when some type of hierarchical routing is used. We examine two types of hierarchical routing schemes and the corresponding end-to-end connection level models. One is fixed or near fixed routing with the typical example being OSPF, which is widely used for Internet, IP based routing. Under this routing scheme, routes are established based on the shortest distance principle, with ties broken according to lower IP address. Considering the fact that links normally fail on a much larger time scale compared to connection durations, this is a fixed routing scheme.

The other type is dynamic/state dependent/adaptive hierarchical routing with the typical example being PNNI. Various proposals for QoS routing in the Internet also fall under this category [8], [9]. In this case, the centering point is “partial information”. Networks are divided into clusters or peer groups. Each non-border nodes is only aware of its own peer group. Routes are established on different layers based on complete information within a peer group and aggregated information between peer groups. The advantage of having a hierarchical end-to-end model is that it closely corresponds to the hierarchical nature of routing and uses only partial information on different layers. By segregating network into layers we can also develop models for situations where different routing schemes are used within a group and between groups.

Substantial numerical experiments are in progress. In the next section we describe network abstraction and aggregation. Hierarchical models for fixed hierarchical routing and dynamic hierarchical routing are presented in Section 3 and 4, respectively. In Section 5 we present preliminary numerical results for the fixed hierarchical routing case, which gained approximately 4-fold improvement in computational cost. Section 5 concludes the paper.

II. NETWORK ABSTRACTION

We only consider large networks that have either physical hierarchies or routing hierarchies vs. a complete mesh since a

Prepared through collaborative participation in the Advanced Telecommunications/Information Distribution Research Program (ATIRP) Consortium sponsored by the U.S. Army Research Laboratory under the Federal Research Laboratory Program, cooperative agreement DAAL01-96-2-0002 and under NASA cooperative agreement NCC3-528.




Fig. 1. Network with three peer groups – Layer One




Fig. 2. Network with three peer groups – Layer Two

hierarchical model promises clear incentives only for the former even if it is at all possible for the latter. Throughout the paper we will be using a two-layer example shown in Figure 1.

There are three peer groups in this example, with the dashed-circles surrounding each one. Each group/node has an address. All border nodes are shown in black. A non-border node does not necessarily have a direct link connected to border nodes, although this is often true with IP networks. Note that all links on this layer are actual, physical links.

All border nodes are kept in the higher layer – in this case Layer 2 and border nodes belonging to the same peer group are fully connected via “logical links”, illustrated in dashed lines if they do not correspond to a physical link as shown in Figure 2.

As pointed out in [10], creating a logical link between each pair of border nodes is the *full-mesh* approach, while collapsing the entire peer group into a single point is the *symmetric-point* approach. Our aggregation approach is a full-mesh one. While it may not be the most economic way of aggregation, this model clearly reflects best the underlying network physical structure and routing scheme. It’s worth pointing out that a bandwidth parameter is usually assigned to a logical link, e.g., representing the maximum/average available bandwidth on the paths between two border nodes, and this may cause problems when different paths overlap [9]. However, as we will see, bandwidth is not necessarily the parameter in our model for calculation on the higher layer, thus avoiding the aforementioned problem. As described in detail in later sections, in the fixed routing case, this parameter is the blocking probability

which resulted from previous iterations within the peer group, and in the dynamic routing case, this parameter can be implied costs, hop number or other criteria based on the dynamic/QoS routing policies being used.

III. HIERARCHICAL MODEL FOR FIXED ROUTING

A. Notations

$G(1.n)$: the n^{th} cluster/peer group on Layer 1, where $n = 1, \dots, N_1$, and N_1 is the total number of peer groups in Layer 1.

$1.n.x$: node x in peer group $G(1.n)$, where $x = 1, \dots, X_n$, and X_n is the total number of nodes in $G(1.n)$.

$1.n.y$: border nodes in peer group $G(1.n)$.

$1.n.x_1 \rightarrow 1.n.x_2$: link from node $1.n.x_1$ to node $1.n.x_2$. Links in our model are directional.

$\lambda_s(1.n_1.x_1 \rightarrow 1.n_2.x_2)$: offered load for class- s traffic from source $1.n_1.x_1$ to destination $1.n_2.x_2$, where $s = 1, \dots, S$, and S is the total number of different traffic classes. It is also written as λ_{ps} with p as the p^{th} source-destination node pair.

$\mathcal{P}(1.n_1.x_1 \rightarrow 1.n_2.x_2)$: the route set between node $1.n_1.x_1$ and $1.n_2.x_2$. \mathcal{P}_p is the route set for the p^{th} node pair.

B. Route and Route Segments

For our modeling purposes, each route is broken down into route segments whenever a route goes across the border. Therefore, a typical route $\mathcal{P}(1.n_1.x_1 \rightarrow 1.n_2.x_2)$ is segmented into the following k segments, assuming that $n_1 \neq n_2$ and that neither $(1.n_1.x_1)$ nor $(1.n_2.x_2)$ is a border node:

$$\begin{aligned} \mathcal{P}^1(1.n_1.x_1 \rightarrow 1.n_1.y_{12}) & \quad \mathcal{P}^2(1.n_1.y_{12} \rightarrow 1.n_i.y_{i1}) \\ \mathcal{P}^3(1.n_i.y_{i1} \rightarrow 1.n_i.y_{i2}) \dots & \quad \mathcal{P}^k(1.n_2.y_{21} \rightarrow 1.n_2.x_2) \end{aligned}$$

where the subscript in y_{j1} indicates this is a border node from which traffic comes into peer group j , and y_{j2} indicates this is a border node from which traffic leaves peer group j . We denote the set of route segments for the p^{th} source-destination node by \mathcal{P}_p .

C. Initial Offered Load and Local Relaxation

The offered load of class- s traffic of the p^{th} node pair $(1.n_1.x_1, 1.n_2.x_2)$ is $\lambda_s^0(1.n_1.x_1 \rightarrow 1.n_2.x_2)$. We substitute this with a combination of the following, in a similar way as route segmentation:

$$\begin{aligned} \lambda_{ps}^0(1.n_1.x_1 \rightarrow 1.n_1.y_{12}) & \quad \text{src. peer grp.} \\ \lambda_{ps}^0(1.n_1.y_{12} \rightarrow 1.n_i.y_{i1}) & \quad \text{inter-peer grp.} \\ \lambda_{ps}^0(1.n_i.y_{i1} \rightarrow 1.n_i.y_{i2}) & \quad \text{intermediate peer grp. } i \\ & \quad \dots \\ \lambda_{ps}^0(1.n_2.y_{21} \rightarrow 1.n_2.x_2) & \quad \text{des. peer grp.} \end{aligned}$$

These terms all take on the value of the real offered load $\lambda_s^0(1.n_1.x_1 \rightarrow 1.n_2.x_2)$. Thus we have complete input information (together with route segments) for each peer group.

For the i^{th} peer group, offered loads indexed with same node pairs are added up to represent the aggregated traffic. Here we assume that at least one of the nodes is a border node since no such additional process is necessary with the case where both nodes are non-border nodes within the same group. Without loss of generality, assume that the destination node is a border node.

$$\lambda_s^1(1.n_i.x_1 \longrightarrow 1.n_i.y_{i2}) = \sum_{\{p: \mathcal{P}(1.n_i.x_1 \longrightarrow 1.n_i.y_{i2}) \in \mathcal{P}_p\}} \lambda_{ps}^0(1.n_i.x_1 \longrightarrow 1.n_i.y_{i2}).$$

The fixed point method is then applied to every peer group separately using these offered loads to calculate group-wide end-to-end blocking probabilities: $B_s(1.n_1.x_1 \longrightarrow 1.n_1.y_{12})$, $B_s(1.n_i.y_{i1} \longrightarrow 1.n_i.y_{i2})$, $B_s(1.n_2.y_{21} \longrightarrow 1.n_2.x_2)$. By doing so, the initial condition of the algorithm is chosen to be of zero inter-group blocking.

D. Reduced Load and Higher Layer Relaxation

On the higher layer (second layer in our example), only border nodes exist. We construct a new network with border nodes, inter-group links and logical links as illustrated in Figure 2. For this logical network we have the following offered load:

$$\lambda_s^1(1.n_1.y_{12} \longrightarrow 1.n_2.y_{21}) = \lambda_s^0(1.n_1.y_{12} \longrightarrow 1.n_2.y_{21}) + \sum_{\{p: \mathcal{P}(1.n_1.y_{12} \longrightarrow 1.n_2.y_{21}) \in \mathcal{P}_p\}} \lambda_{ps}^0(1.n_1.y_{12} \longrightarrow 1.n_2.y_{21}) \cdot B_s(1.n_1.x_1 \longrightarrow 1.n_1.y_{12}) \cdot B_s(1.n_2.y_{21} \longrightarrow 1.n_2.x_2),$$

This is the initial offered load thinned by blocking in both the source and destination peer groups. For simplicity purposes, we use $\mathcal{P}(1.n_1.y_{12} \longrightarrow 1.n_2.y_{21})$ for either a single route segment or combination of multiple route segments belonging to the same route.

We again apply the fixed point approximation to this layer and calculate second-layer end-to-end blocking probabilities. The result of this step is the group-to-group blocking probability: $B_s(1.n_1.y_{12} \longrightarrow 1.n_2.y_{21})$.

E. Iteration

Using the results from the inter-group approximation, replace the offered load with:

$$\lambda_s^1(1.n_i.x_1 \rightarrow 1.n_i.y_{i2}) = \sum_{\{p: \mathcal{A} \in \mathcal{P}_p\}} \lambda_{ps}^0(1.n_i.x_1 \rightarrow 1.n_i.y_{i2}) \cdot B_s(1.n_i.y_{i2} \rightarrow 1.n_j.y_{j1}) \cdot B_s(1.n_j.y_{j1} \rightarrow 1.n_j.x_2).$$

where \mathcal{A} is defined as the union: $\mathcal{P}(1.n_i.x_1 \rightarrow 1.n_i.y_{i2}) \cup \mathcal{P}(1.n_i.y_{i2} \rightarrow 1.n_j.y_{j1}) \cup \mathcal{P}(1.n_j.y_{j1} \rightarrow 1.n_j.x_2)$. This is essentially the original offered load thinned by blocking on inter-group links and the remote peer group. This becomes the new input for local relaxation. Local and higher layer relaxations are then repeated till the difference between results from successive iterations are within certain criteria.

IV. HIERARCHICAL MODEL FOR DYNAMIC ROUTING

There are numerous existing and proposed dynamic/QoS hierarchical routing schemes, each of which results in different end-to-end performances determined by the scope and design trade-off of the routing scheme. Our primary goal here is not to design an end-to-end model for each and everyone of these schemes. Rather, we attempt to present an end-to-end performance modeling framework that considers a “generic” type of dynamic hierarchical routing, which captures some of the most basic properties of a majority of such routing schemes. We make assumptions for simplicity purposes, but our work shows how an end-to-end performance model can be closely coupled with routing policies to provide an efficient way of analysis. Furthermore, our model enables us to analyze situations where different routing schemes are used on difference levels of a networks.

A. Dynamic Hierarchical Routing

One key property of any dynamic hierarchical routing is *inaccurate/incomplete information*. A node has only aggregated information on other peer groups advertised by the border nodes. This aggregated information can be one or more of various metrics specified by the routing algorithm: implied cost of a peer group maximum available bandwidth between border node pairs, end-to-end blocking or delay incurred by going through a peer group, etc..

In source routing, a path is selected with detailed hop-by-hop information in the originating peer group but only group-to-group information beyond the originating group. The choice of routes within a group can be determined using shortest path routing, least loaded/state dependent routing and so on. The routes between groups are primarily determined by the form of aggregation advertised by border nodes. A call is blocked if the route selected according to the dynamic routing policy does not have the required bandwidth.

B. Probabilistic Offered Load Distribution and Higher Layer Relaxation

One of the main advantages of dynamic routing is “load averaging”, i.e., dynamically distribute traffic flow onto different paths of the network to achieve greater utilization of network resources. We argue that under steady state, a particular traffic flow (defined by class, source-destination node pair) is distributed among all feasible routes, and among multiple border nodes that connect to other peer groups. (This problem does not exist when there is only one border node. Routes are still dynamically chosen, but all routes ultimately go through that single border node.) The fraction of a traffic flow that goes through a certain border node is directly related to the aggregated information/metrics for the group-to-group route the border node sees.

Based on this, for a pair of nodes belonging to different clusters, the feasible route set are divided into three subsets: routes

within the source peer group, routes between groups and routes within the destination peer group.

The offered load for the class- s traffic for node pair $(1.n_1.x_1 \rightarrow 1.n_2.x_2)$ is $\lambda_s^0(1.n_1.x_1 \rightarrow 1.n_2.x_2)$, and each route between peer groups (second-layer route) gets a portion:

$$\begin{aligned}\lambda_{ps}^0(\mathcal{P}^2(1.n_1.y_1 \rightarrow 1.n_2.y_1)) &= a_1 \lambda_s^0(1.n_1.x_1 \rightarrow 1.n_2.x_2) \\ \lambda_{ps}^0(\mathcal{P}^2(1.n_1.y_2 \rightarrow 1.n_2.y_2)) &= a_2 \lambda_s^0(1.n_1.x_1 \rightarrow 1.n_2.x_2)\end{aligned}$$

$$\dots$$

$$\lambda_{ps}^0(\mathcal{P}^2(1.n_1.y_{n_p} \rightarrow 1.n_2.y_{n_p})) = a_{n_p} \lambda_s^0(1.n_1.x_1 \rightarrow 1.n_2.x_2)$$

where $a_i, i = 1, 2, \dots, n_p$ is the fraction of traffic going through each of the valid route set. $\sum_i a_i = 1$.

For simplicity purposes, denote $2.y_i$ as any node on the second layer.

So the aggregated traffic for node pairs on the second layer is

$$\lambda_s^1(2.y_i \rightarrow 2.y_j) = \sum_{\{p: \mathcal{P}(2.y_1 \rightarrow 2.y_2) = \mathcal{P}_p^2\}} \lambda_{ps}^0(\mathcal{P}^2).$$

We thus have all the input traffic load for the second layer and the fixed point method for a flat network with dynamic/state dependent routing can be applied [7]. This results in the end-to-end blocking probability $B_s(2.y_i \rightarrow 2.y_j)$.

As discussed earlier, different criteria(delay, blocking probability, implied cost, available bandwidth) associated with the second segments of the same original traffic flow should match the distribution of traffic flow onto these segments. Ultimately one of the goals for any dynamic routing scheme is to balance traffic load on different alternative routes, and the end result is that these alternative routes will have equivalent QoS under steady state. For example, we can use blocking probability as a criteria to adjust the traffic distribution $a_i, i = 1, 2, \dots, n_p$. Segments with a blocking probability higher than median gets a decreased portion, and segments with a blocking probability lower than median gets an increased portion:

$$\begin{aligned}a_i &:= a_i + \delta \quad \text{if } B_s(1.n_1.y_i \rightarrow 1.n_2.y_i) < B_m; \\ a_i &:= a_i - \delta \quad \text{if } B_s(1.n_1.y_i \rightarrow 1.n_2.y_i) > B_m,\end{aligned}$$

where δ is a small incremental value and B_m is the median blocking probability among all routes. Other means of relating traffic distribution to route QoS can also be specified. Another round of iteration is then started using these new distribution values. This process continues until all routes have similar blocking probabilities.

C. Lower Layer Relaxation

From the offered load distribution calculated from the higher layer relaxation, we now have complete traffic information for each peer group, including the traffic when the group is a source group, a destination group or an intermediate group.

The reduced load, which is the above thinned by blocking on the second layer and the remote peer group, becomes the input offered load for calculations on this layer in the same way to that in the fixed routing model.

Iteration is done in a similar way to that with higher layer. This will result in a new set of values of traffic distribution, which is then used by the next iteration on the second layer.

V. NUMERICAL AND SIMULATION RESULTS

In this section we present numerical experiment results for the network example shown in Fig 1 using fixed hierarchical routing scheme, and dynamic hierarchical routing scheme. This is a 21-node, 30-link, 3-clusters, 2-layer network model. We use single class of traffic requiring unit bandwidth. Link capacities varies from 80 to 160. We use the offered traffic load between node pairs at a ‘‘nominal’’ level. The intensity of this traffic load is shown in Table I, in which *load* is defined as the ratio between the total rate of traffic coming out of a node and the total out-going link bandwidth connecting to this node. At the nominal level, the value of this ratio for each node is around 0.05. In addition to this offered traffic load we also define a ‘‘weight’’ in our experiment as a multiplier to the nominal traffic, so that we get twice, three times of the nominal traffic, etc.. The complete data on link capacities and traffic rates can be found in [11].

A. Fixed Hierarchical Routing

Since when using fixed hierarchical routing a network can always be treated as flat, we compare the performance of flat fixed-point approximation (FPA) and the hierarchical fixed-point approximation. It can be shown [1] that under certain limiting regime for fixed routing the fixed point approximation is asymptotically correct.

Table II is a comparison between flat fixed-point approximation and hierarchical fixed-point approximation on individual link blocking probabilities (end-to-end blocking probabilities

TABLE I
NOMINAL OFFERED TRAFFIC LOAD

Node	Rate	Cap.	load
1.1.1	11.15	270	0.041296
1.1.2	7.20	160	0.045000
1.1.3	10.60	180	0.058889
1.1.4	7.90	150	0.052667
1.1.5	8.60	210	0.040952
1.1.6	8.95	220	0.040682
1.1.7	10.45	180	0.058056
1.2.1	8.95	210	0.042619
1.2.2	8.80	220	0.040000
1.2.3	6.65	130	0.051154
1.2.4	9.15	180	0.050833
1.2.5	9.70	180	0.053889
1.2.6	7.95	140	0.056786
1.3.1	9.55	230	0.041522
1.3.2	12.0	290	0.041379
1.3.3	3.70	80	0.046250
1.3.4	7.90	140	0.056429
1.3.5	8.25	160	0.051562
1.3.6	9.00	180	0.050000
1.3.7	8.00	130	0.061538
1.3.8	5.75	100	0.057500

are computed directly from these for fixed routing). We used seven times nominal traffic (weight = 7).

TABLE II
COMPARISON OF RESULTS, WEIGHT = 7.

link	Hier. FPA	Flat FPA
(1.1.1-1.1.5)	0.235701	0.235704
(1.1.1-1.2.1)	0.409961	0.409955
(1.1.2-1.2.6)	0.526158	0.526159
(1.1.4-1.1.6)	0.148341	0.148333
(1.1.5-1.1.6)	0.006189	0.006190
(1.1.2-1.1.7)	0.000000	0.000000
(1.1.5-1.3.1)	0.004523	0.004523
(1.1.6-1.3.2)	0.054767	0.054766
(1.2.2-1.3.1)	0.122970	0.122967
(1.2.2-1.2.4)	0.000007	0.000007
(1.3.3-1.3.5)	0.000000	0.000000
time (sec)	13.90	42.76

We see that the hierarchical scheme gives very close results compared to that of the flat approximation scheme, but achieved 3 ~ 4-fold improvement in computation.

B. Dynamic Hierarchical Routing

We use the same network example, with shortest path routing within each cluster, but use least-loaded routing between clusters. Least-loaded routing (LLR) is a form of bandwidth-optimization QoS routing [12]. A source node chooses a border node based on the advertised average blocking between the border node and the destination cluster, and a border node chooses the route that has the most free capacity among all route from itself to the destination cluster. We used the reduced load model for least-loaded routing we developed in [7]. The distribution of traffic among border nodes is inversely proportional to the advertised blocking probability.

The following tables show the comparison between the results of the hierarchical model and the discrete event simulation (DES), with weight being 5, 10, 15 and 20, respectively.

We see as the traffic increases, the model generates better approximations. Overall the approximation is satisfactory. The run time for the approximation is around 11-15 seconds while the simulation typically takes 5-20 minutes to converge depending on the traffic load.

TABLE III
COMPARISON OF RESULTS, WEIGHT = 5.

Node Pair	Hier. FPA	DES
(1.1.3-1.3.3)	0.00077	0.00000
(1.1.2-1.2.4)	0.00000	0.00000
(1.1.6-1.3.7)	0.00000	0.00000
(1.1.1-1.2.3)	0.00000	0.00000
(1.2.1-1.2.6)	0.00191	0.00453
(1.3.1-1.3.8)	0.00000	0.00000
(1.3.5-1.3.6)	0.00000	0.00000

TABLE IV
COMPARISON OF RESULTS, WEIGHT = 10.

Node Pair	Hier. FPA	DES
(1.1.3-1.3.3)	0.28784	0.25235
(1.1.2-1.2.4)	0.00470	0.00166
(1.1.6-1.3.7)	0.09745	0.09995
(1.1.1-1.2.3)	0.10185	0.11089
(1.2.1-1.2.6)	0.35567	0.36216
(1.3.1-1.3.8)	0.06322	0.05275
(1.3.5-1.3.6)	0.00000	0.000000

TABLE V

COMPARISON OF RESULTS, WEIGHT = 15.

Node Pair	Hier. FPA	DES
(1.1.3-1.3.3)	0.54327	0.52136
(1.1.2-1.2.4)	0.05856	0.04963
(1.1.6-1.3.7)	0.29617	0.30721
(1.1.1-1.2.3)	0.23567	0.24354
(1.2.1-1.2.6)	0.53049	0.54061
(1.3.1-1.3.8)	0.26342	0.25285
(1.3.5-1.3.6)	0.00000	0.00000

TABLE VI

COMPARISON OF RESULTS, WEIGHT = 20.

Node Pair	Hier. FPA	DES
(1.1.3-1.3.3)	0.66527	0.66421
(1.1.2-1.2.4)	0.16374	0.16717
(1.1.6-1.3.7)	0.42122	0.43112
(1.1.1-1.2.3)	0.33676	0.34187
(1.2.1-1.2.6)	0.62975	0.62494
(1.3.1-1.3.8)	0.39753	0.39476
(1.3.5-1.3.6)	0.00000	0.00000

VI. CONCLUSION

In this paper we presented a hierarchical reduced load approximation scheme for networks with either fixed hierarchical routing or dynamic hierarchical routing policies. This is a novel approximation method for efficient and scalable performance analysis. It can also be used in cases where different routing schemes are used in different regions of a network. Our preliminary numerical experiment results showed significant improvement in computational cost. Further simulation and study in algorithm convergence and asymptotic correctness is needed.

REFERENCES

- [1] F. P. Kelly. Loss Networks. *The Annals of Applied Probability*, 1(3):319–378, 1991.
- [2] F. P. Kelly. Notes on effective bandwidth. In *Stochastic Networks Theory and Applications*, F. P. Kelly, S. Zachary and I. Ziedins, pages 141–168. Eds. Oxford, U. K.: Science, 1996.
- [3] D. Mitra, R. Gibbens, and B. D. Huang. State-Dependent Routing on Symmetric Loss Networks with Trunk Reservations – I. *IEEE Trans. Communications*, 41(2):400–411, 1993.
- [4] S. Chung, A. Kashper, and K. W. Ross. Computing Approximate Blocking Probabilities for Large Loss Networks with State-Dependent Routing. *IEEE Trans. Networking*, 1(1):105–115, 1993.
- [5] A. G. Greenberg and R. Srikant. Computational Techniques for Accurate Performance Evaluation of Multirate, Multihop Communication Networks. *IEEE/ACM Trans. Networking*, 5(2):266–277, 1997.
- [6] M. Liu, J. S. Baras, and A. Misra. Performance evaluation in multi-rate multi-hop communication network with adaptive routing. *Proc. Annual ATIRP Conference*, 1:136–140, February 1998.
- [7] M. Liu and J. S. Baras. Fixed point approximation for multirate multihop loss networks with adaptive routing. Technical report, Institute for Systems Research, University of Maryland, College Park, MD, 1999. Technical Report TR99-44.
- [8] J. Behrens and J. J. Garcia-Luna-Aceves. Hierarchical Routing Using Link Vectors. *Proc. INFOCOM*, 1998.
- [9] M. Montgomery and G. de Veciana. Hierarchical Source Routing Through Clouds. *Proc. INFOCOM*, 1998.
- [10] W. C. Lee. Spanning tree method for link state aggregation in large communication networks. *Proc. IEEE Infocom*, 1:297–302, 1995.
- [11] Mingyan Liu. Network Performance Modeling, Design and Dimensioning Methodologies. Technical Report PhD 2000-6, Center for Satellite and Hybrid Communication Networks, University of Maryland, College Park, 2000.
- [12] S. Chen and K. Nahrstedt. An Overview of Quality of Service Routing for Next-Generation High-Speed Networks: Problems and Solutions. *IEEE Network*, 1998.